Data-Driven Control of Cellular Networks

Sandeep Chinchali, Marco Pavone, Sachin Katti

Stanford University

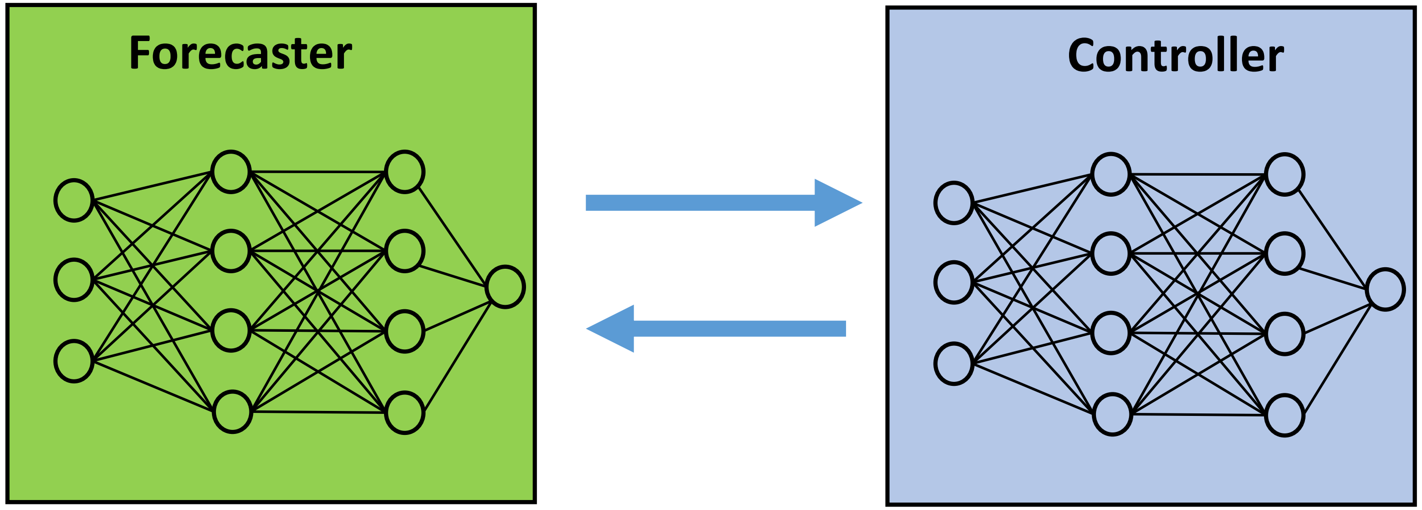

Stanford University

Data-Driven Network Control is Ubiquitous

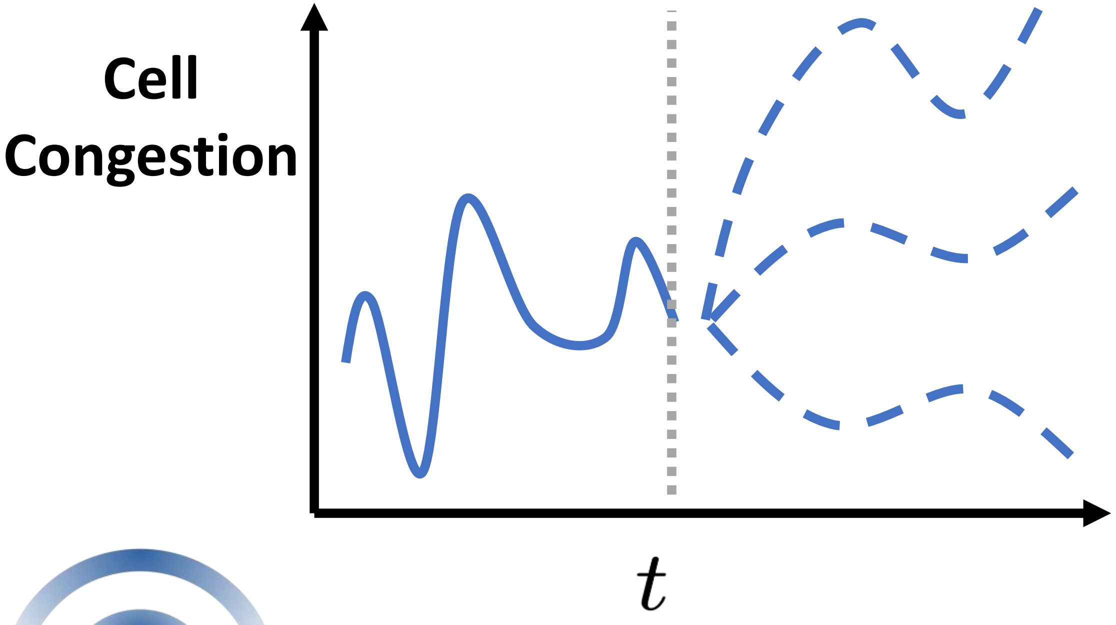

Video Streaming

Robotic Taxi Fleets

IoT Sensor Updates

Optimal Control

Challenges of Network Control

1. Data-driven forecasts
   - What ***features/statistics*** are needed for control?
2. Many Input Variables
   - Forecaster and Controller
3. Increasingly:
   - Data boundaries

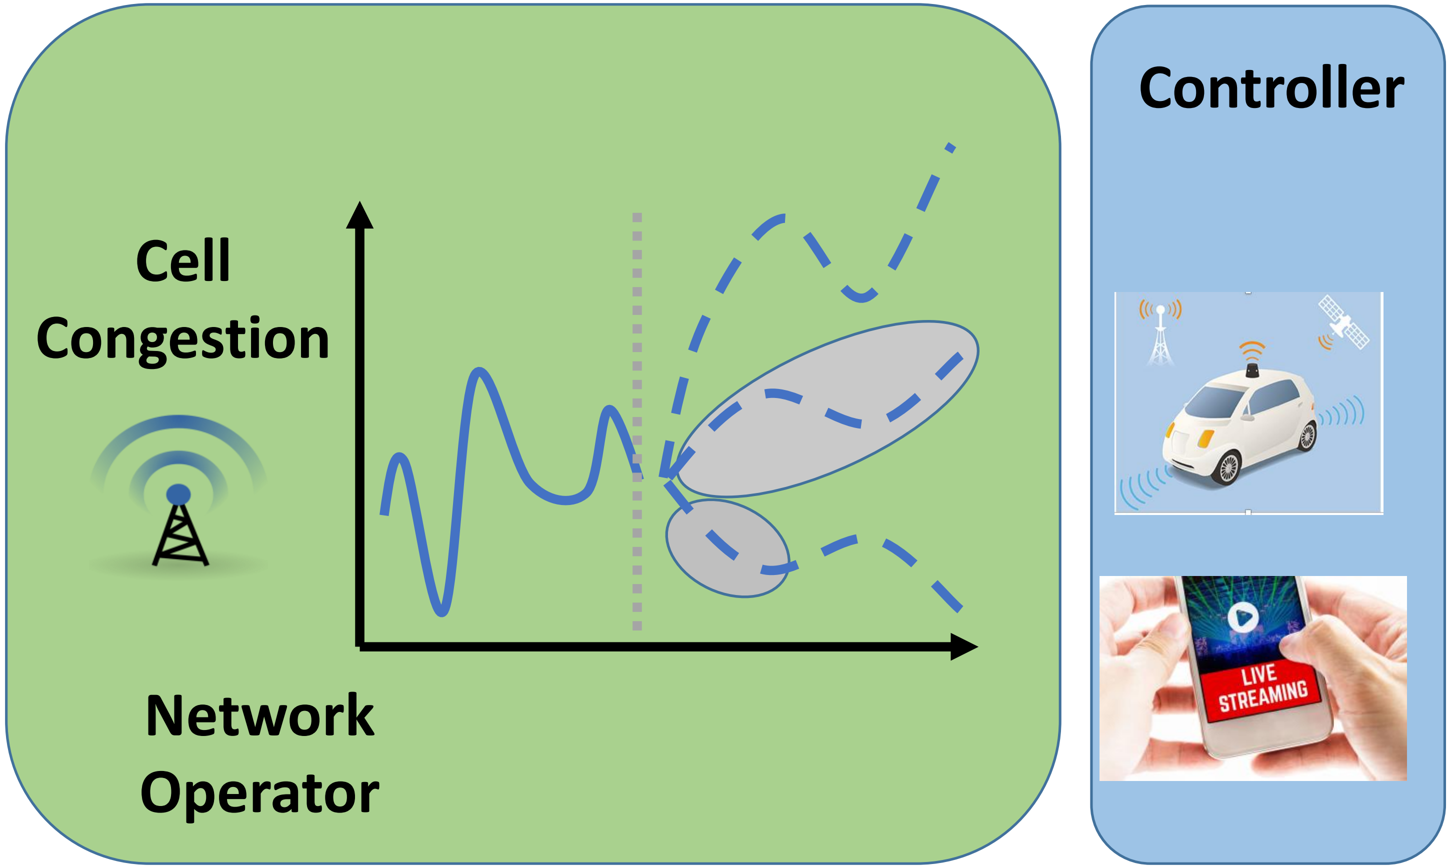

General Approach

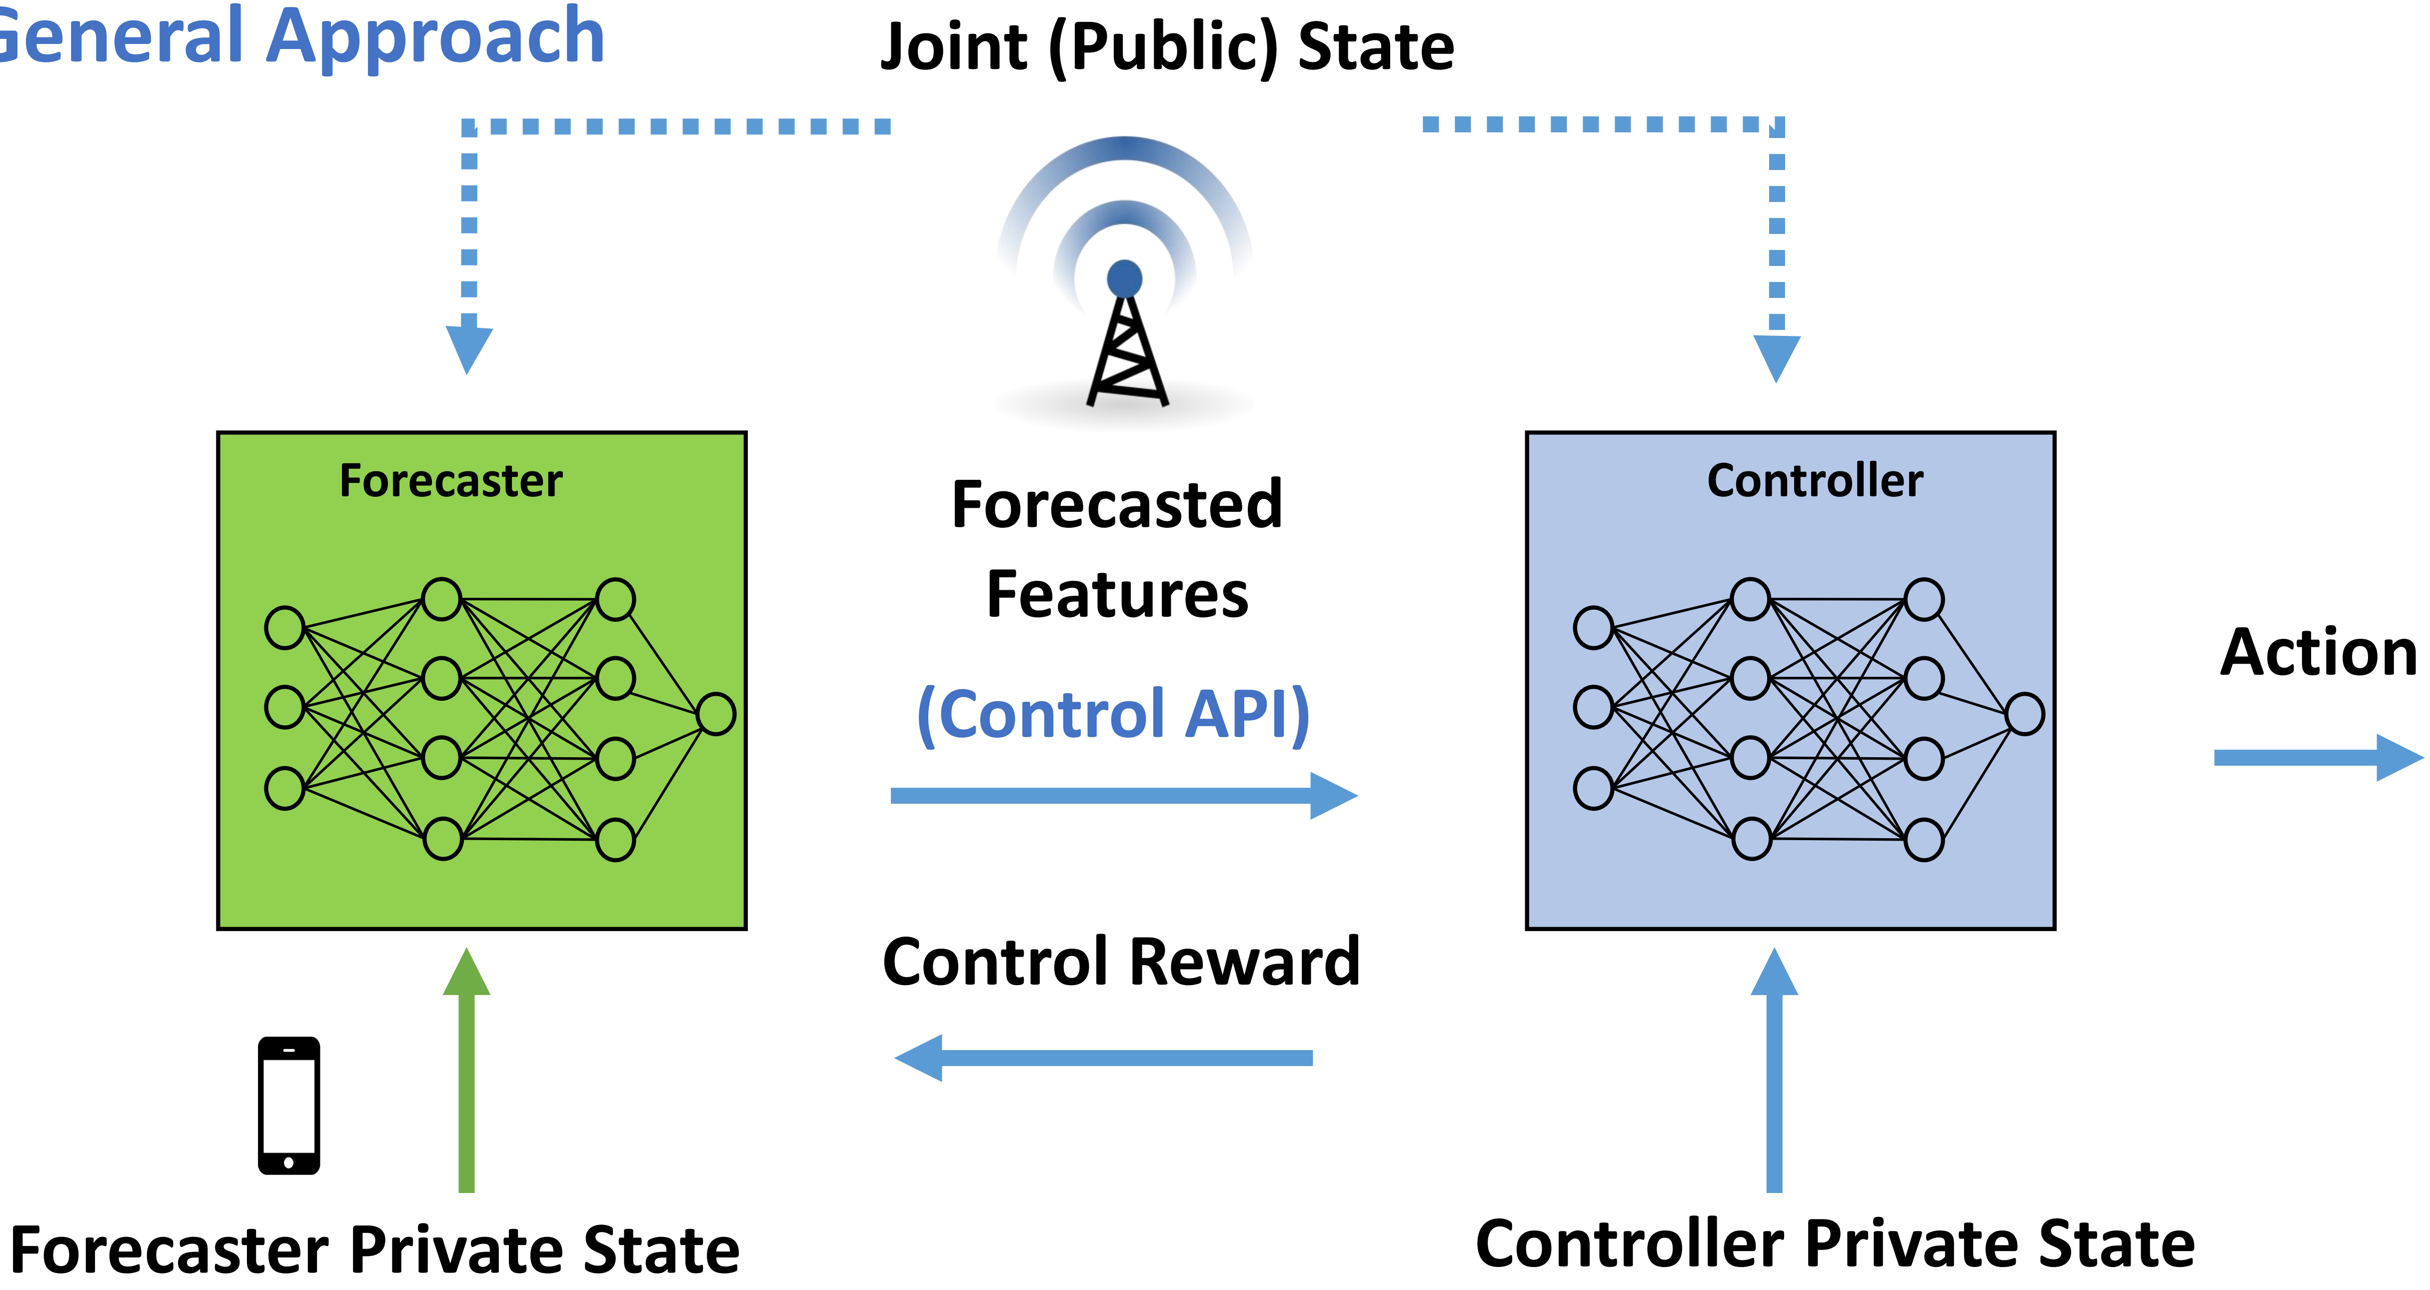

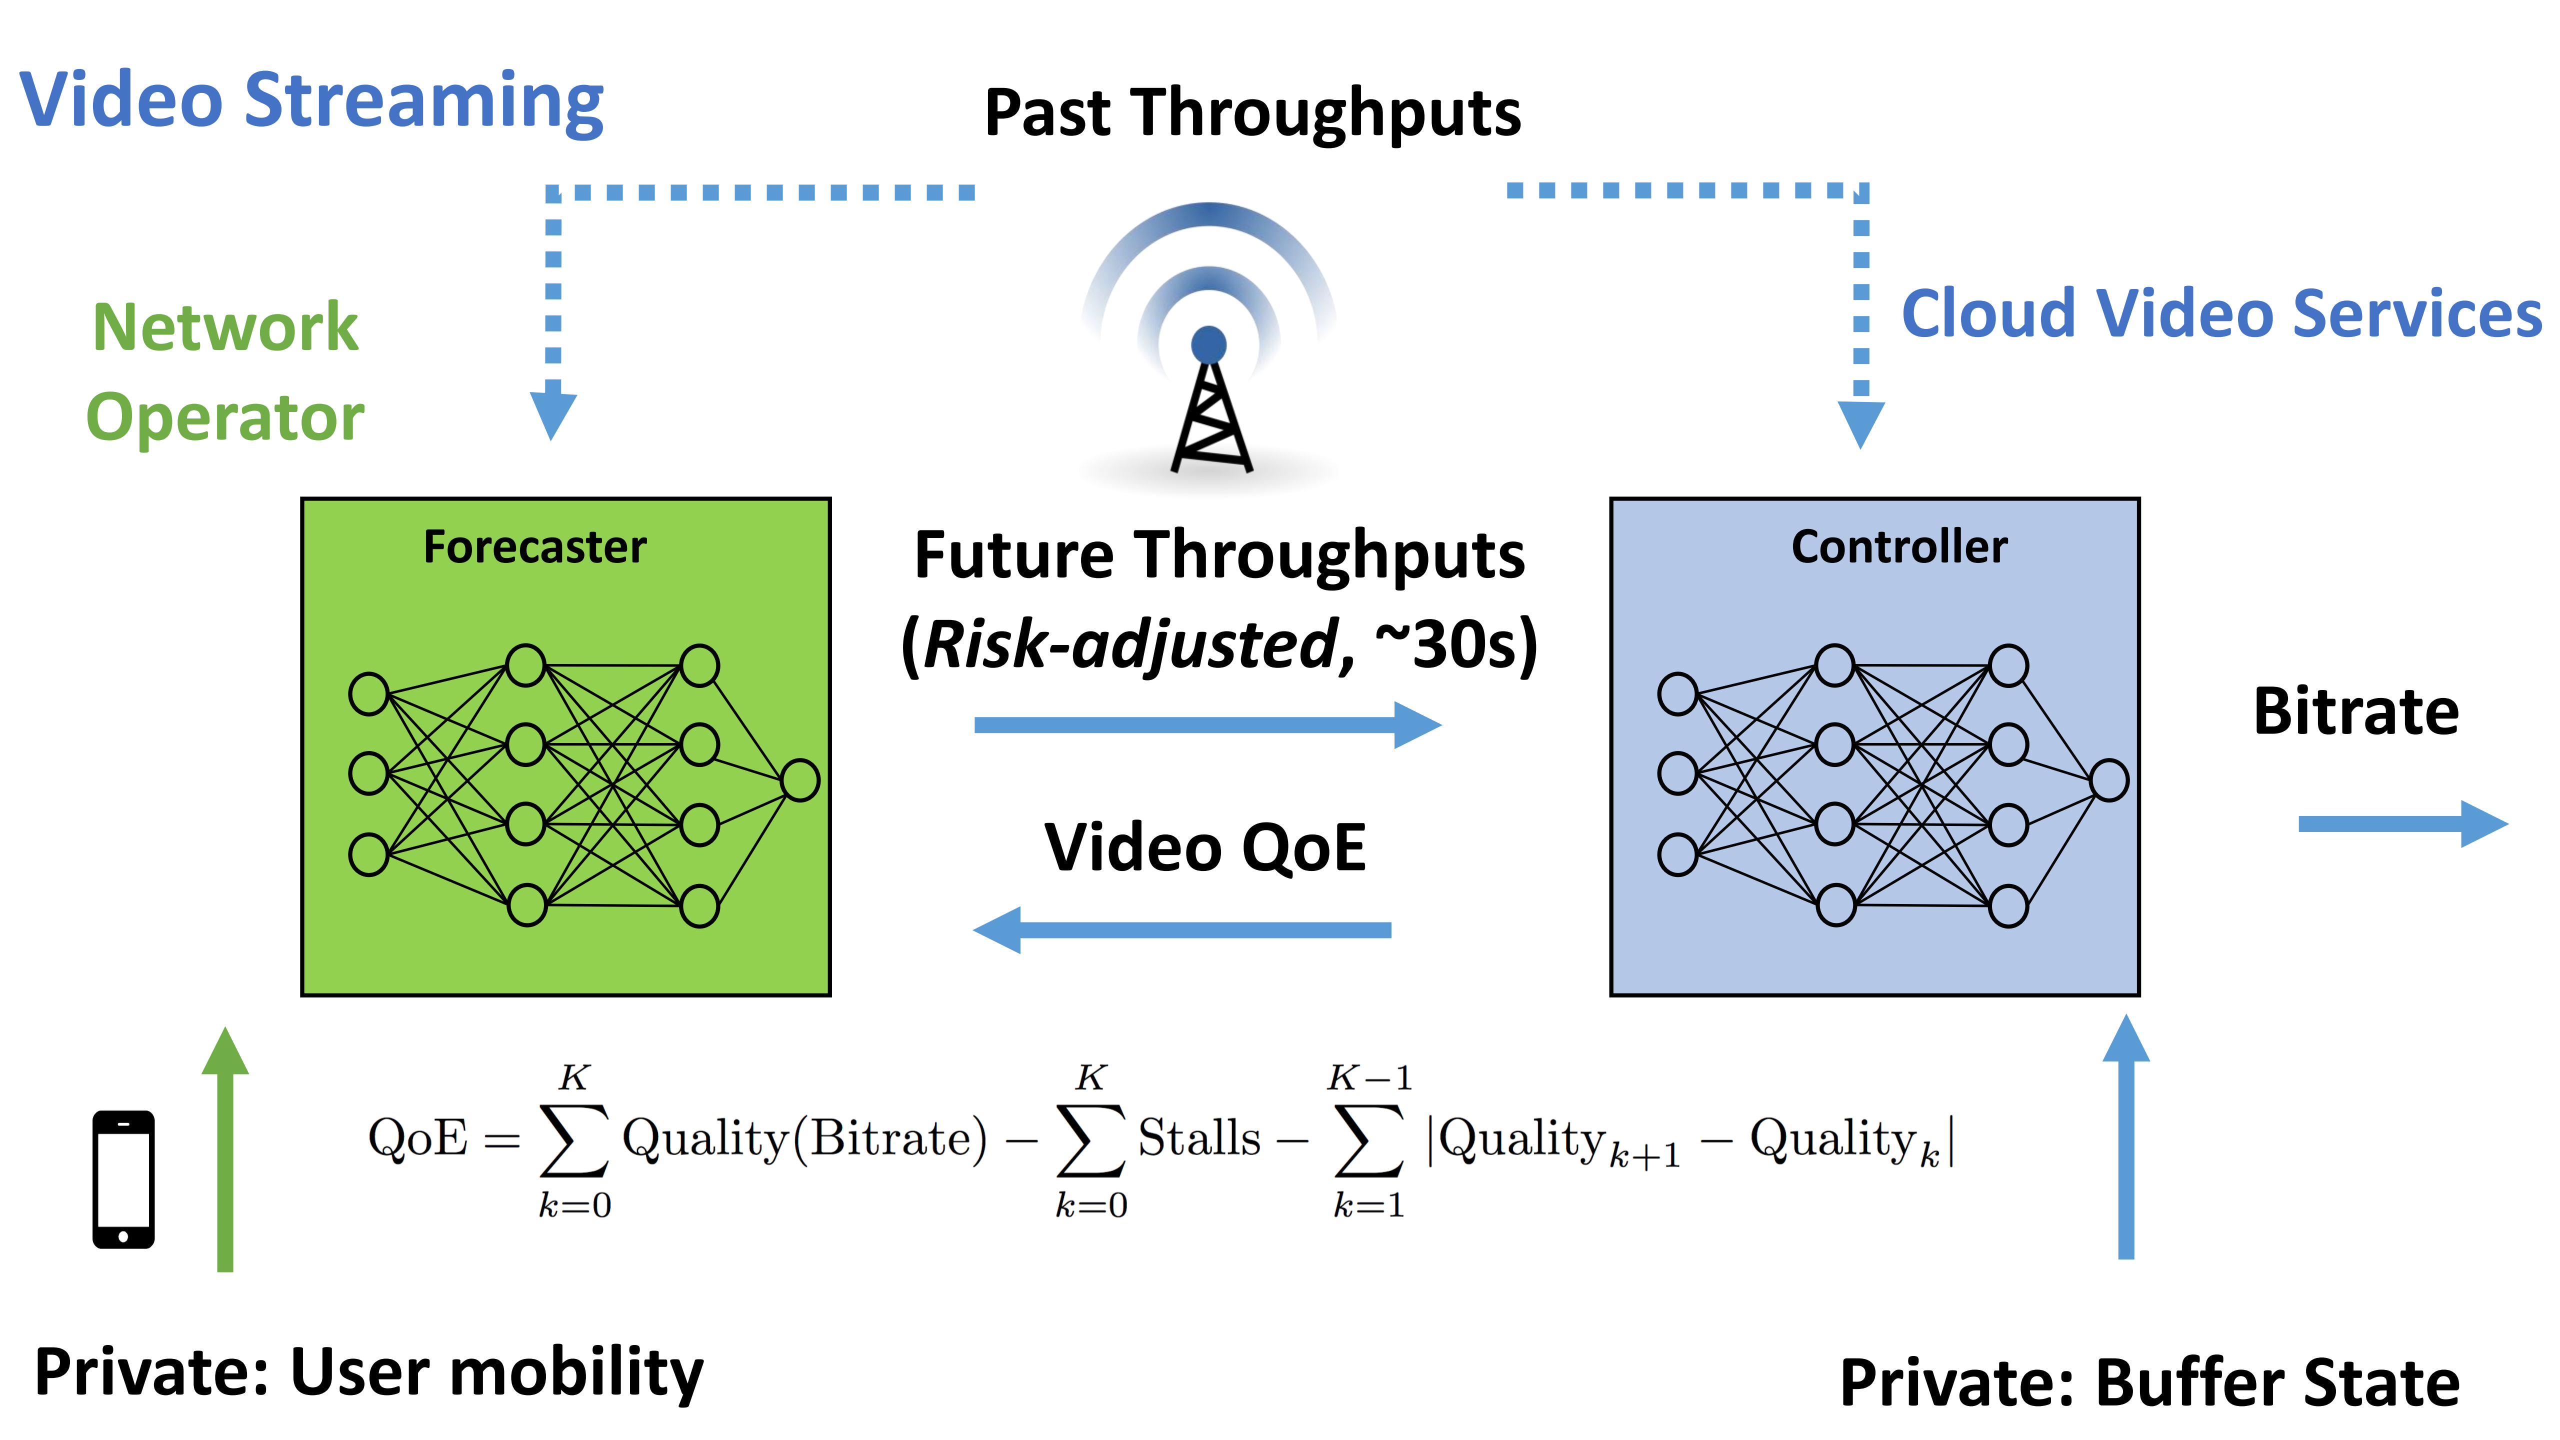
Video Streaming
Past Throughputs
Network
Operator
Cloud Video Services
Forecaster
Future Throughputs
(*Risk-adjusted*, ~30s)
Controller
Bitrate
Video QoE
$\text{QoE} = \sum_{k=0}^{K} \text{Quality}(\text{Bitrate}) - \sum_{k=0}^{K} \text{Stalls} - \sum_{k=1}^{K-1} |\text{Quality}_{k+1} - \text{Quality}_{k}|$
Private: User mobility
Private: Buffer State

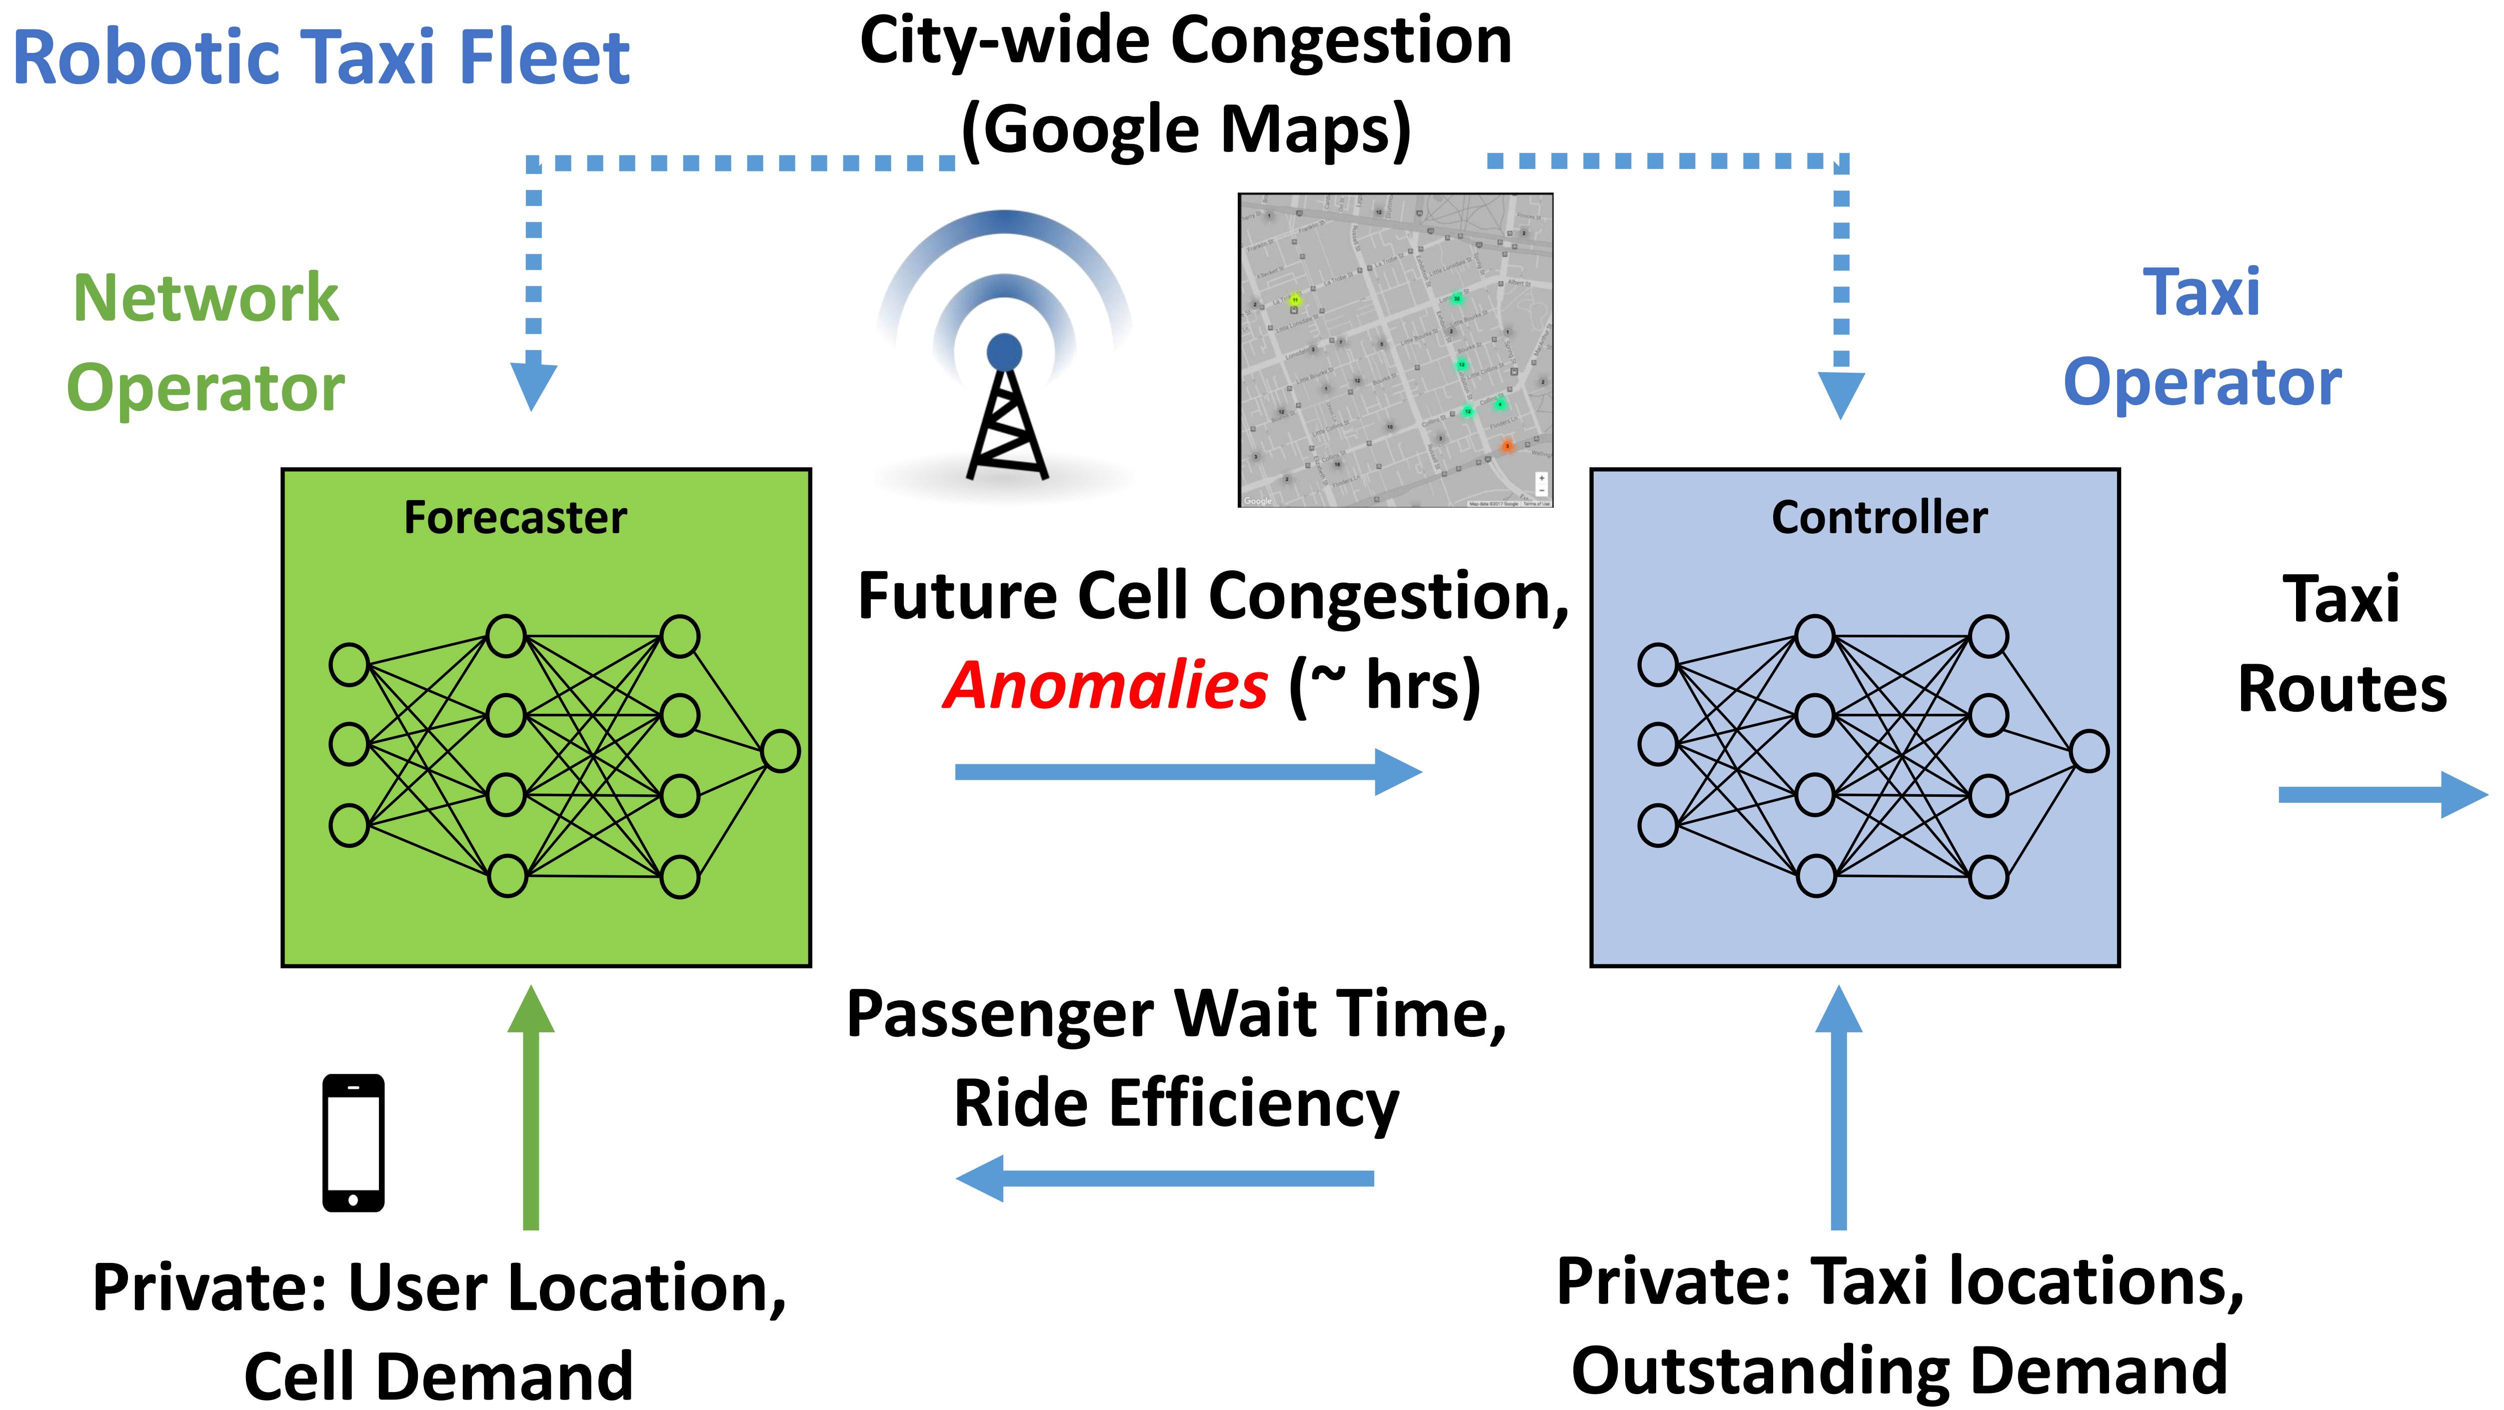
Robotic Taxi Fleet
City-wide Congestion
(Google Maps)
Network
Operator
Taxi
Operator
Forecaster
Controller
Future Cell Congestion,
Anomalies (~ hrs)
Taxi
Routes
Passenger Wait Time,
Ride Efficiency
Private: User Location,
Cell Demand
Private: Taxi locations,
Outstanding Demand

Approach:
Reinforcement Learning (RL)

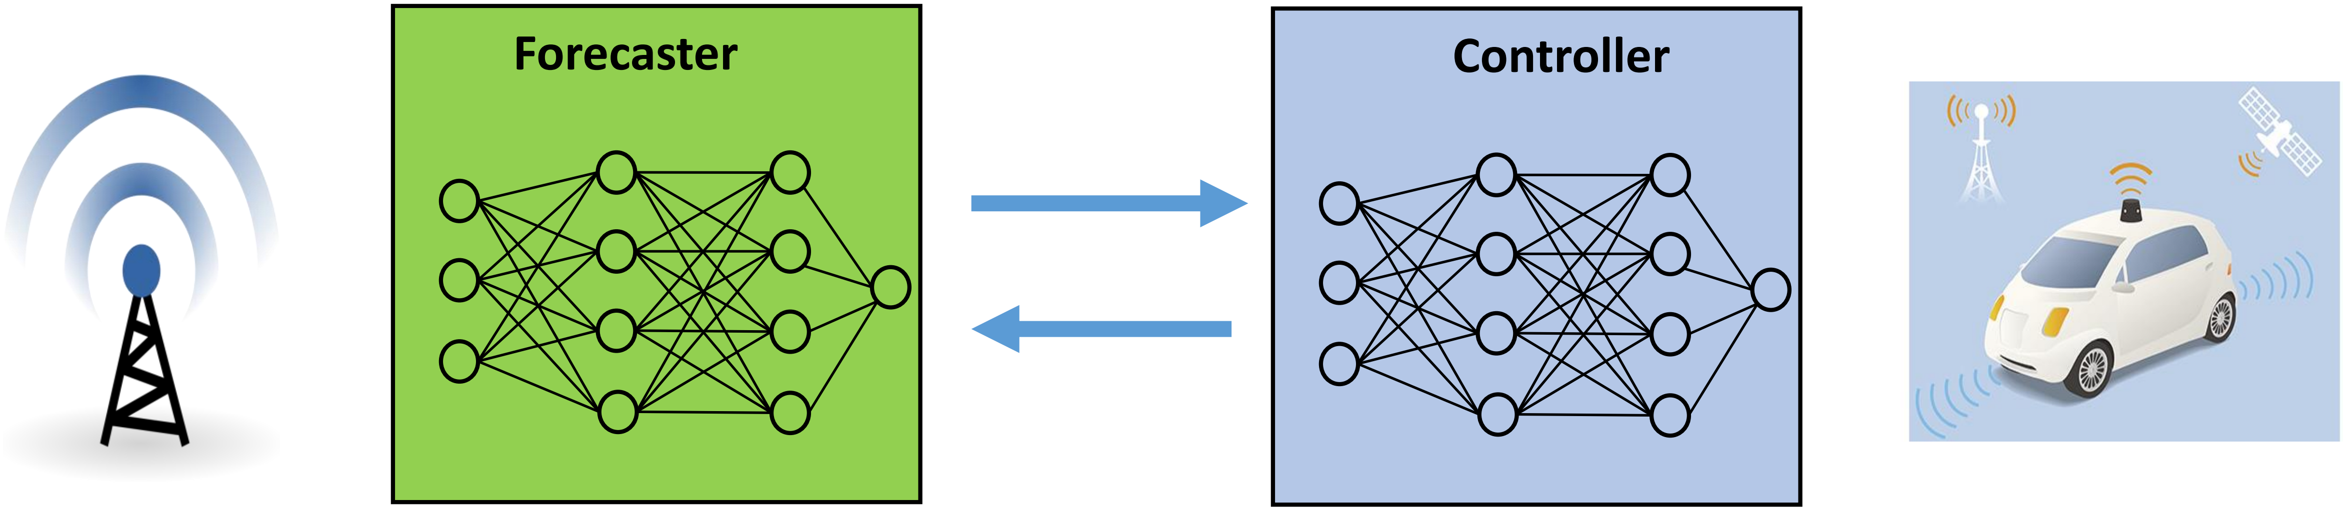

Reinforcement Learning (RL)

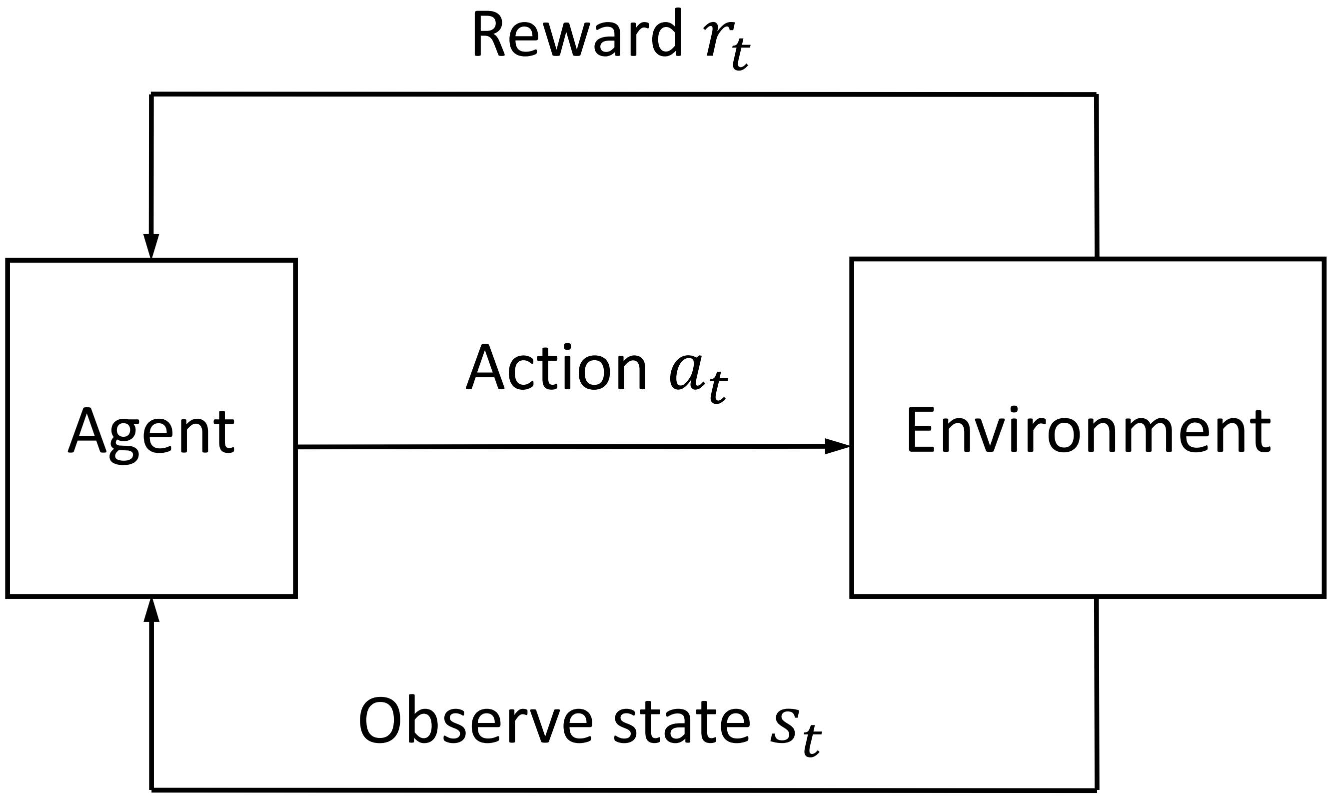

Goal: Maximize the total reward $\sum_t r_t$

Adapted from Pensieve (Sigcomm 18, Mao et. al.)

Cellular Network Traffic Scheduling (AAAI 2018)

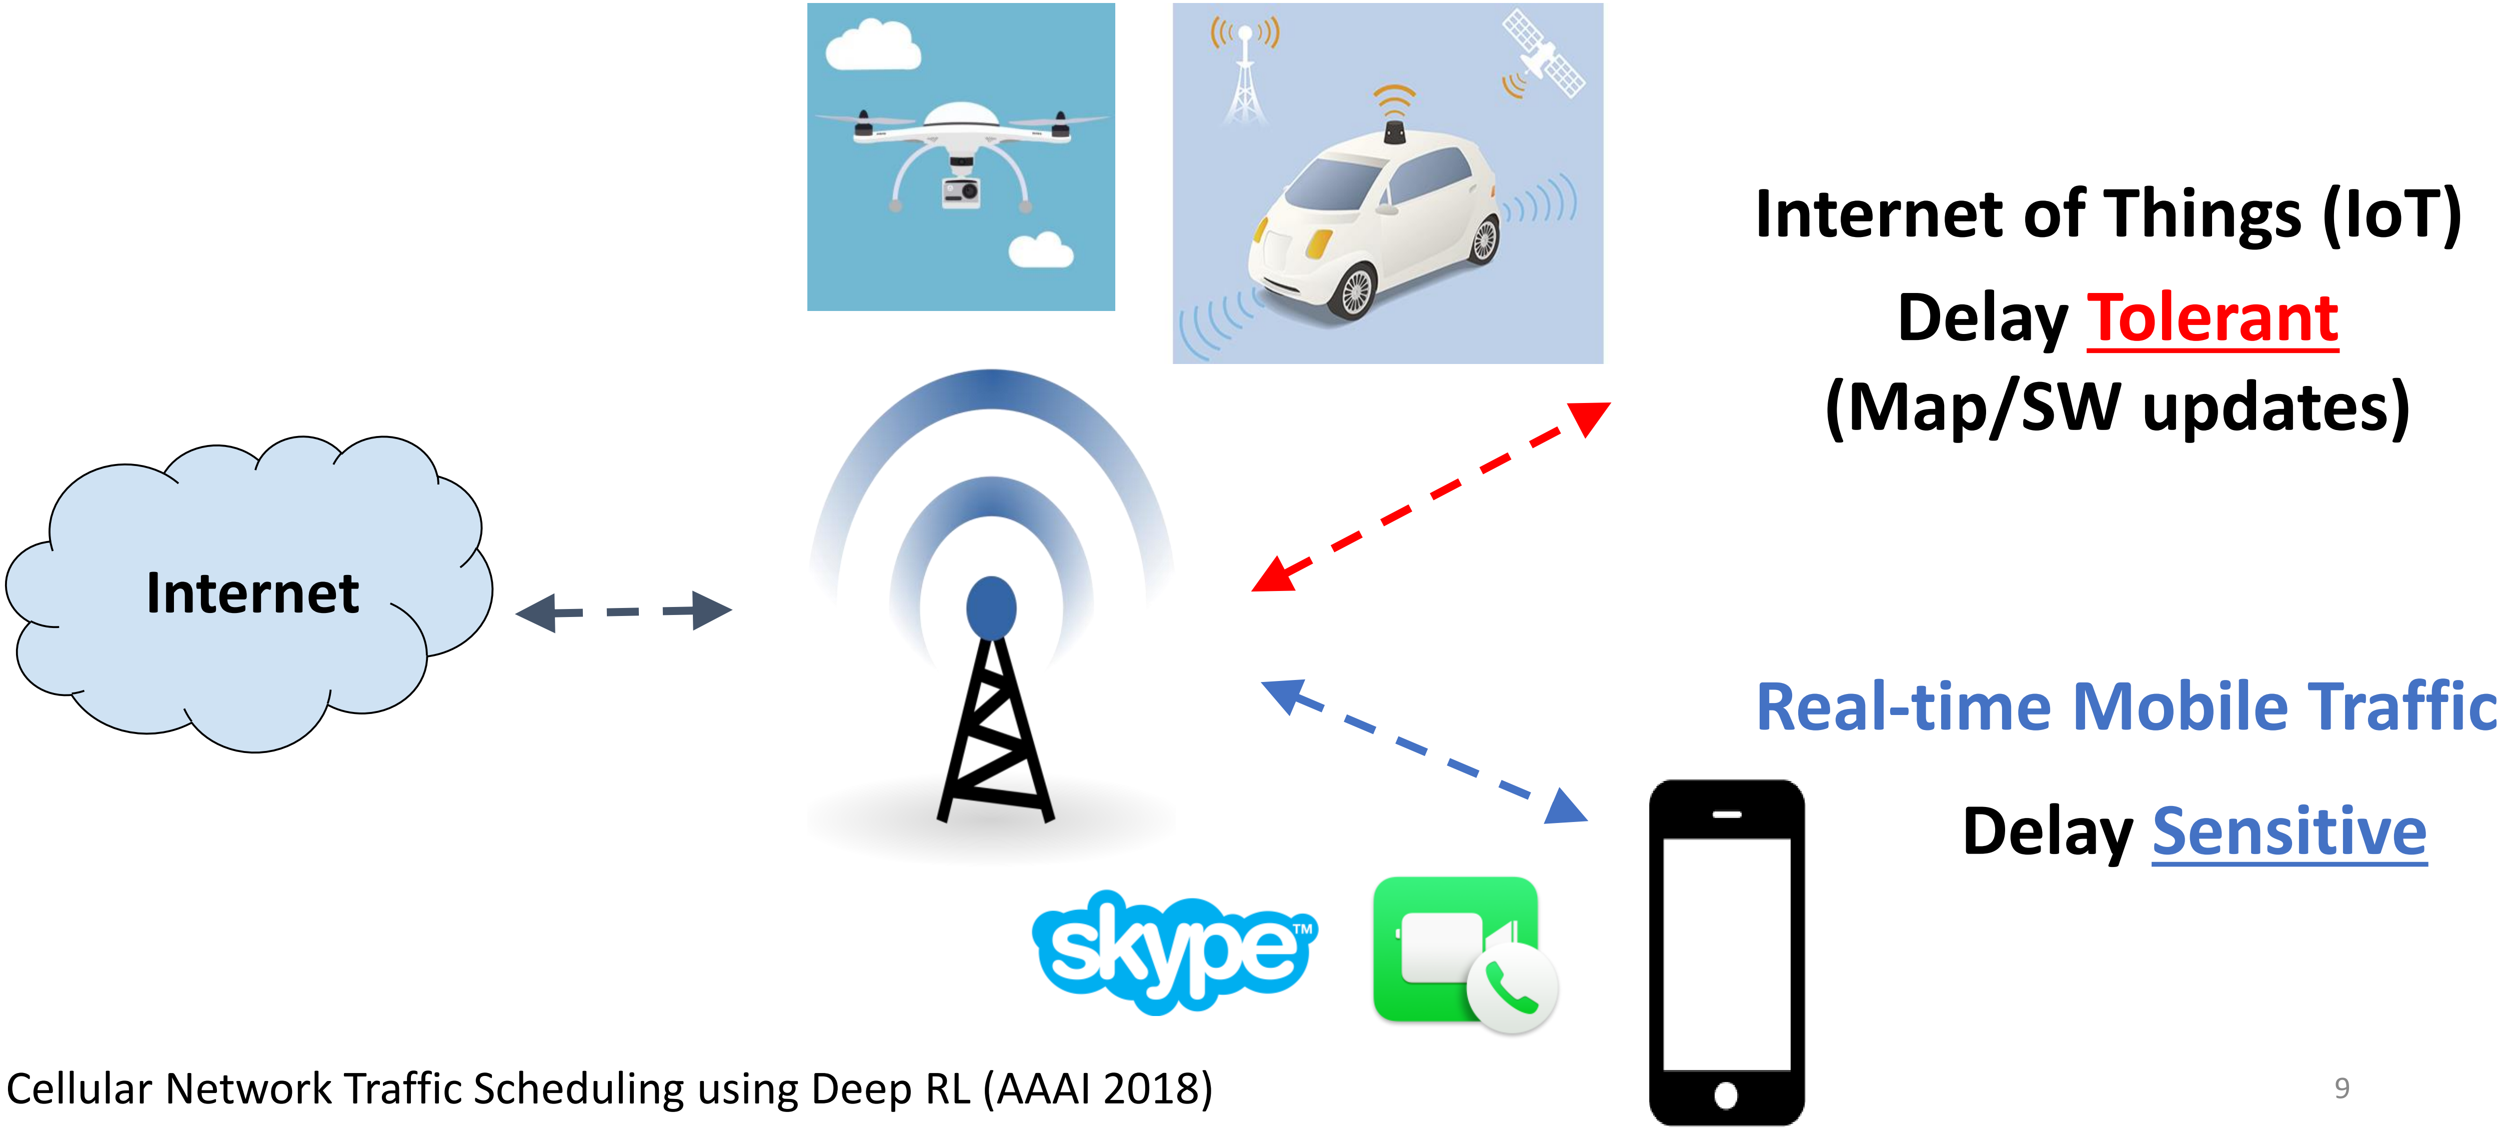

Why is IoT traffic scheduling hard?

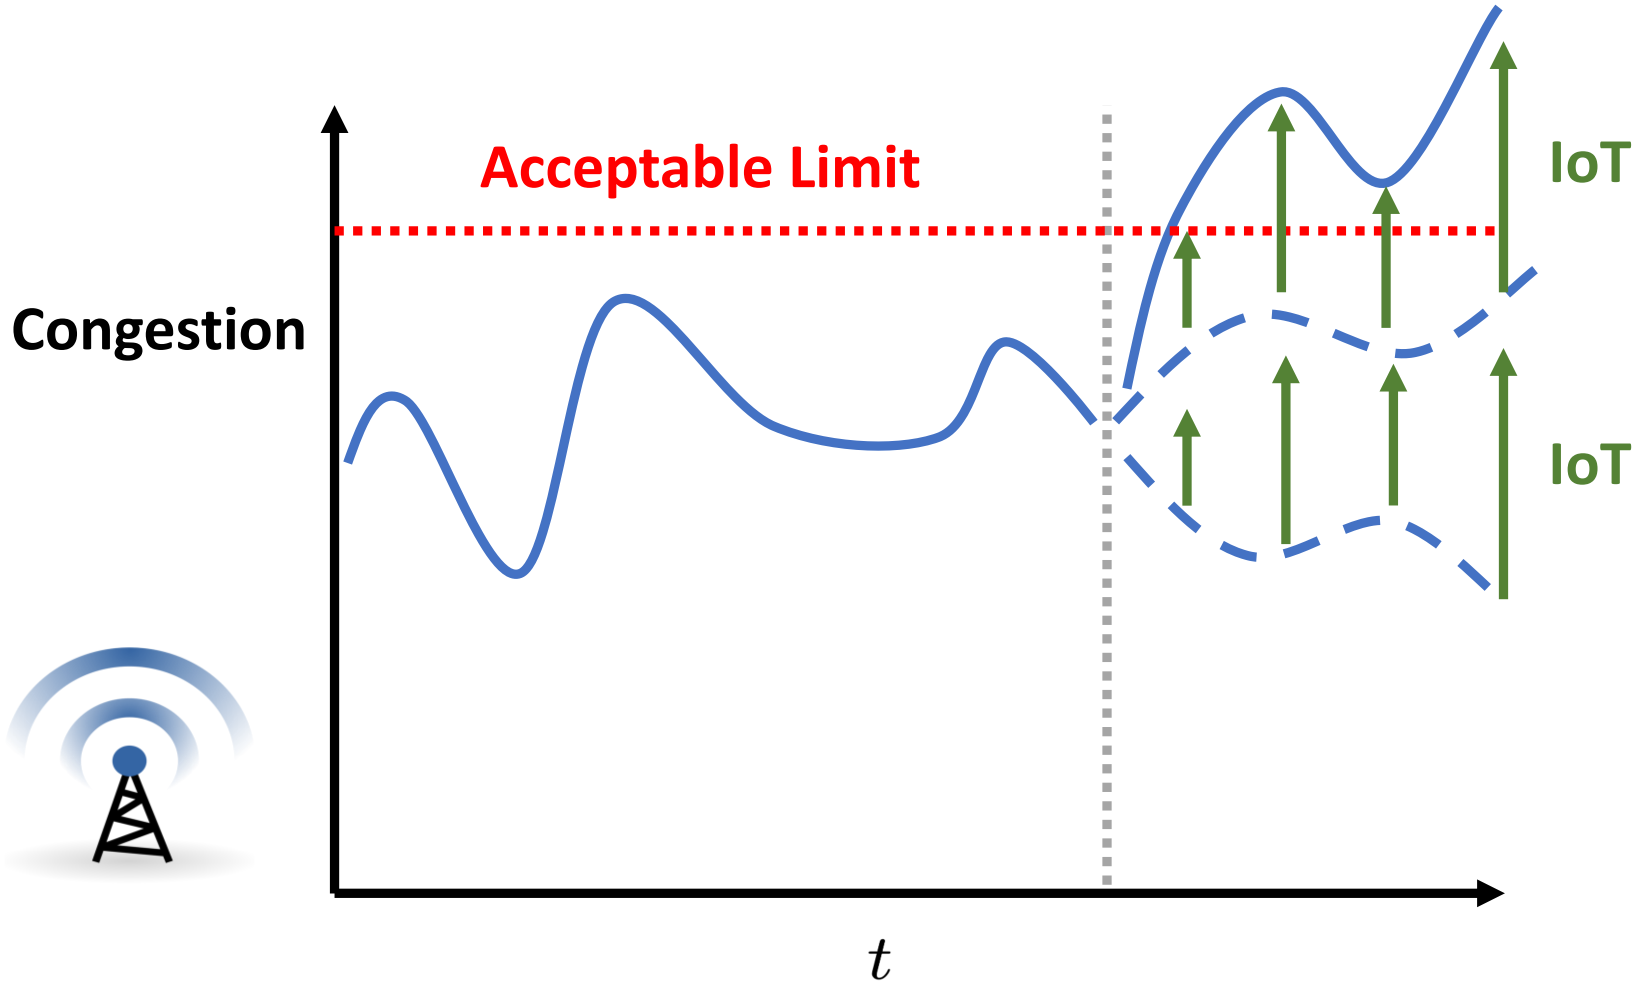

Contending goals

- Max IoT data
- Loss to mobile traffic
- Network limits

Optimal Control

RL Schedules Sensor Updates

1. Network State Space (Cell congestion forecasts)
2. IoT Scheduler Actions (Traffic Rate)
3. Operator policies/reward: efficient use of network

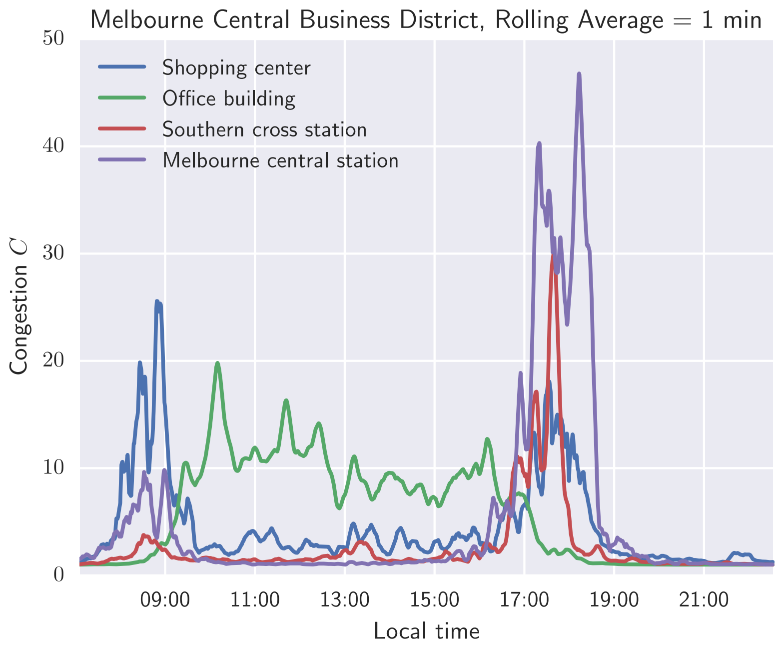

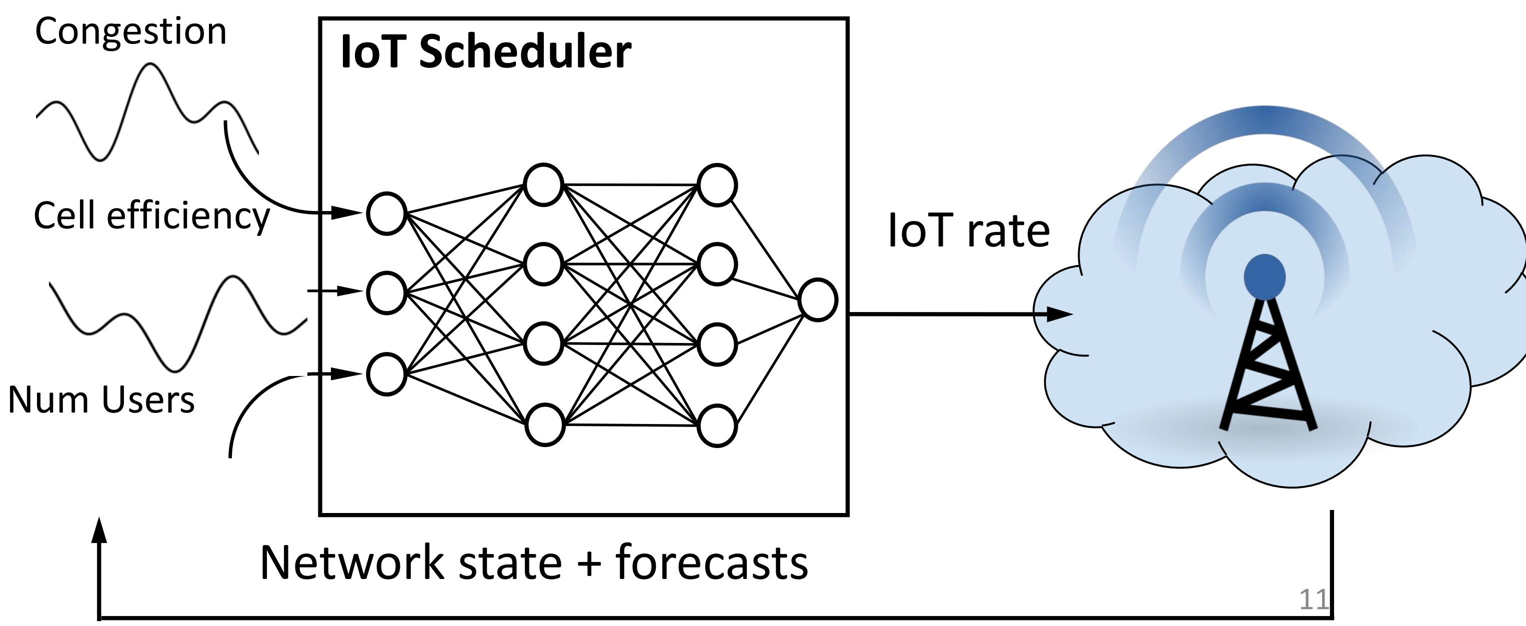

RL Dynamics: Live Network Experiments

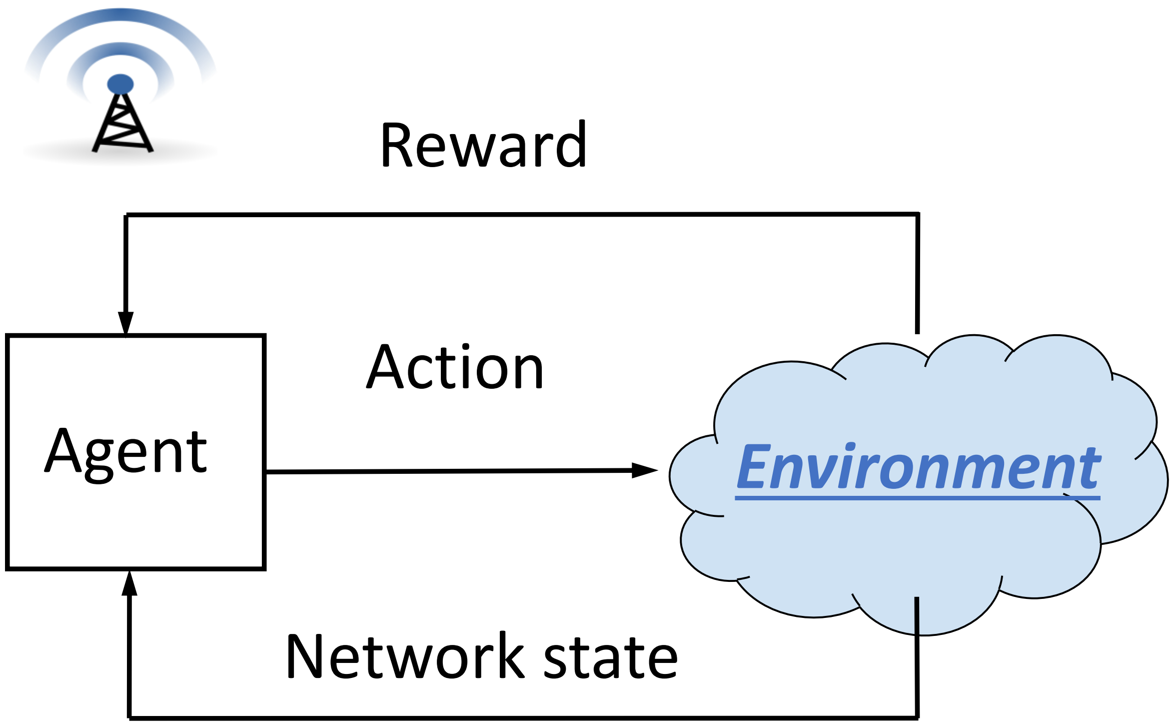

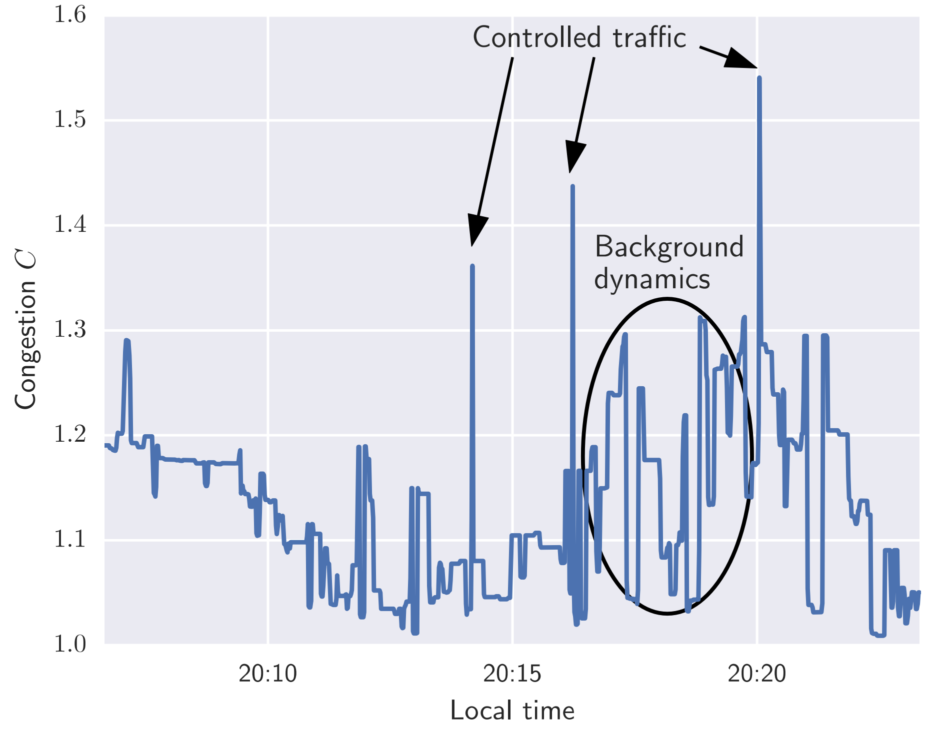

$$p(s_{t+1} | s_t, a_t)$$

$$C_{t+1} = \left\{ \begin{array}{ll} \underbrace{C_t + Ma_t}_{\text{controlled state}} + \underbrace{\Delta\tilde{C}_t}_{\text{historical commute}} + \epsilon_t & \text{if } a_t > 0 \\ \underbrace{\tilde{C}_t + \Delta\tilde{C}_t}_{\tilde{C}_{t+1}} + \epsilon_t & \text{if } a_t = 0 \end{array} \right\}$$

RL exploits transient dips in utilization

Controlled Congestion

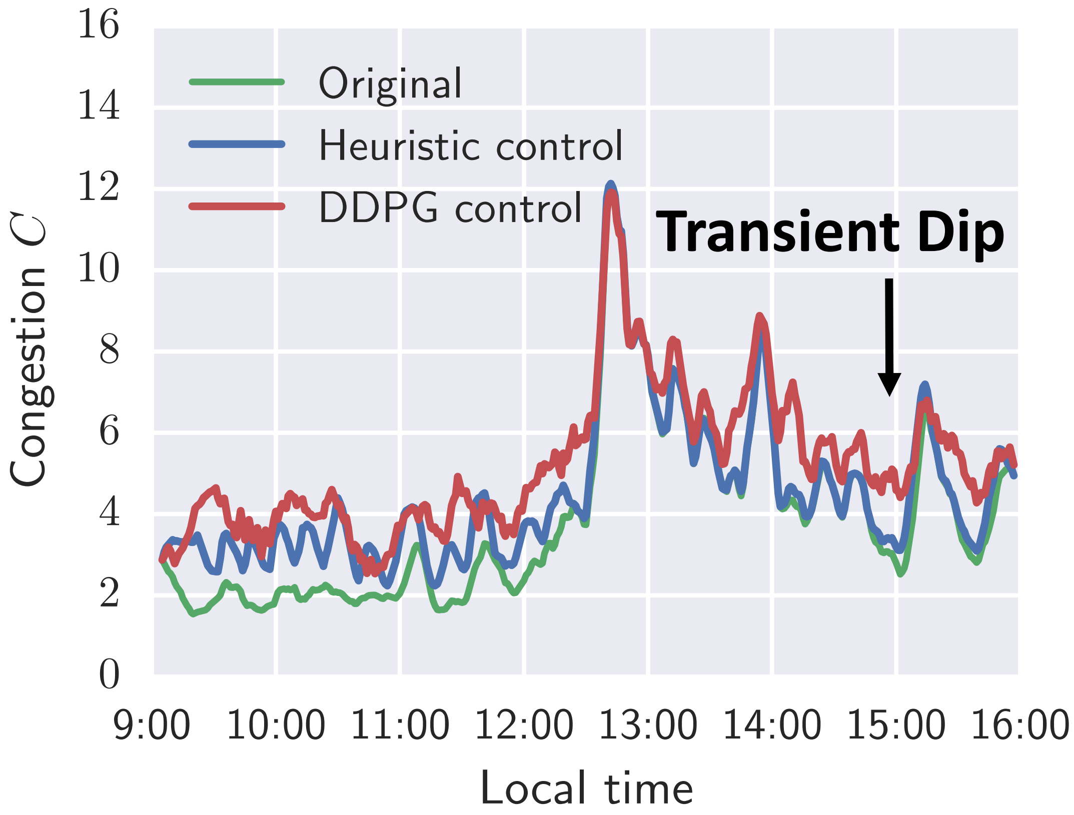

Utilization gain

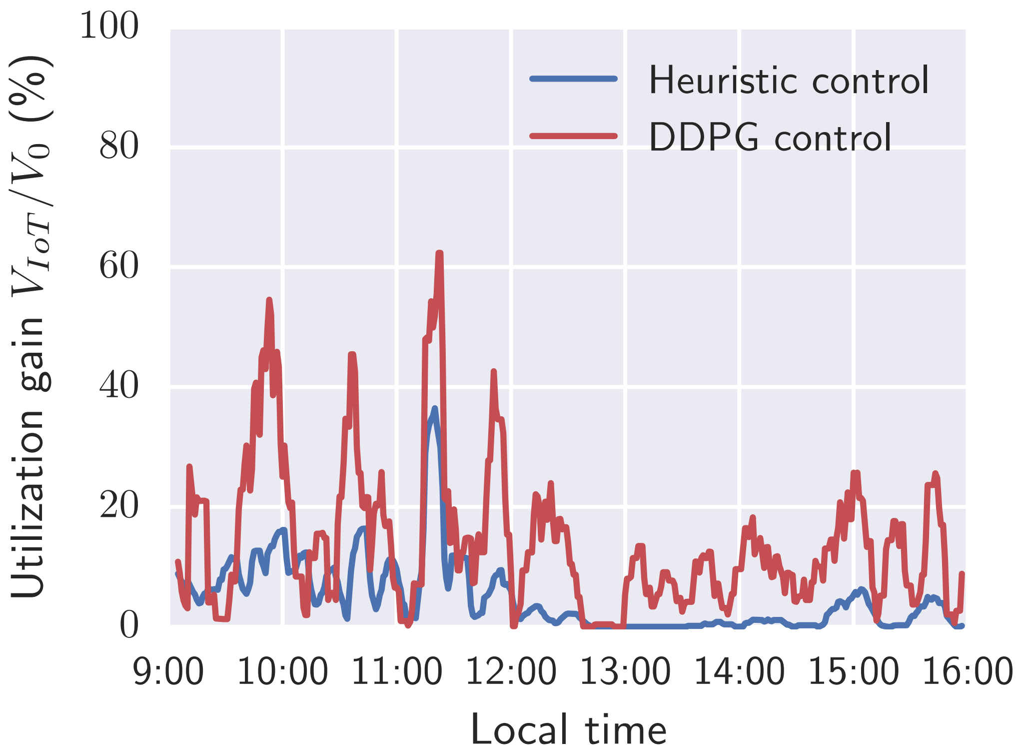

Application 2: Mobile Video Streaming

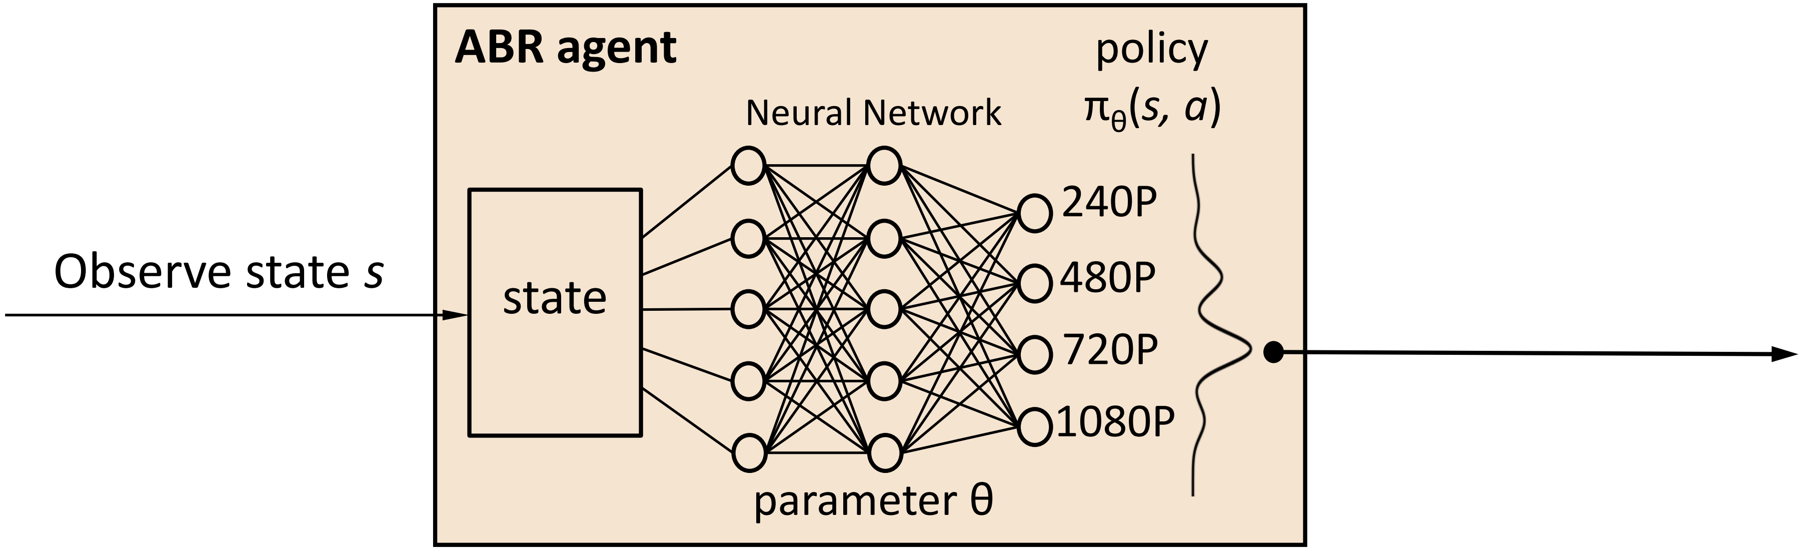

How will forecasts of network conditions improve ABR?

Figure from Pensieve (Sigcomm 18, Mao et. al.)

Palo Alto Cell Throughput Diversity

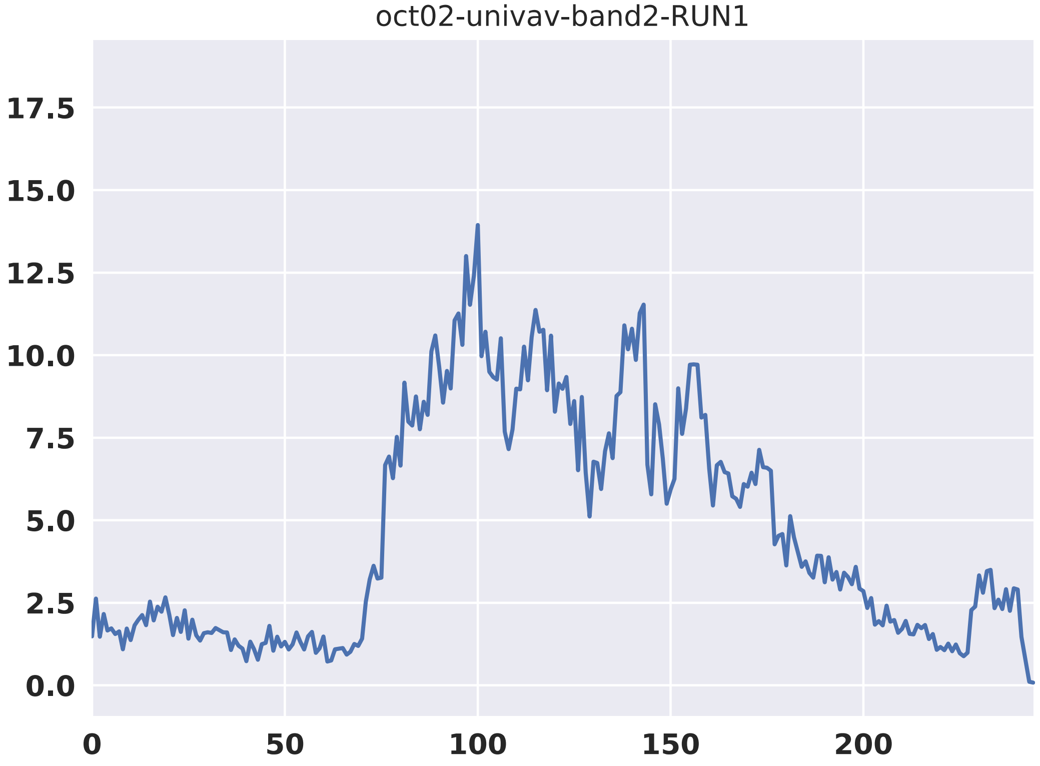

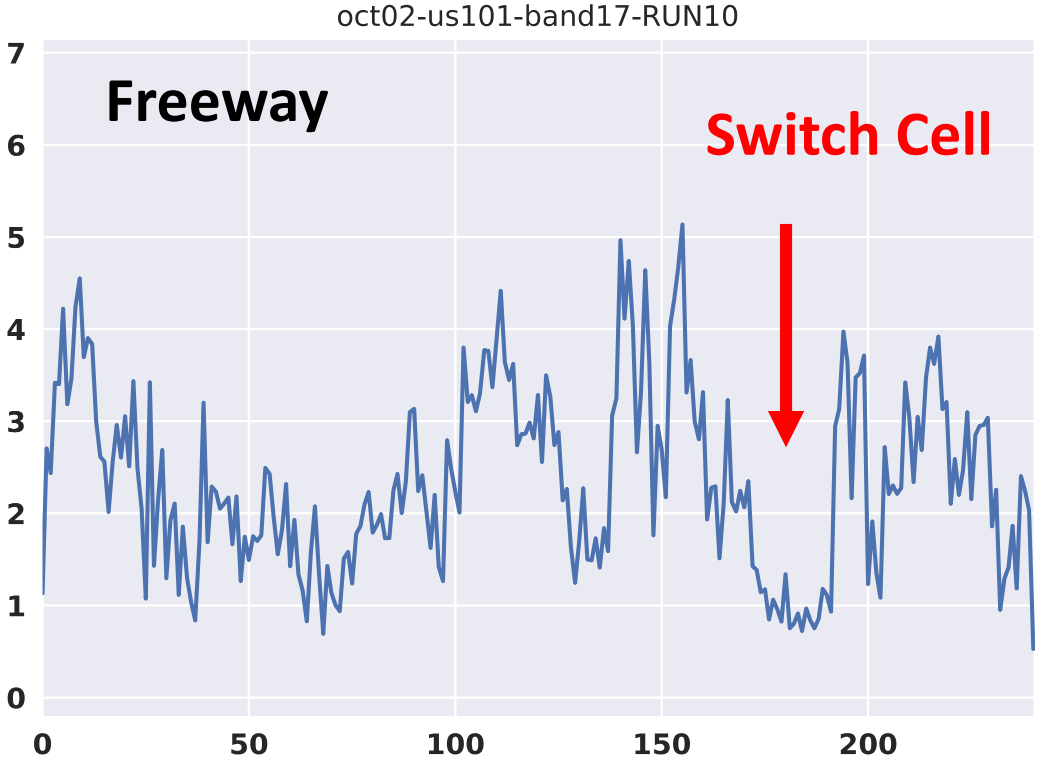

<u>Insight:</u> **Foresight of <u>*true* network condition</u> helps**

<u>Solution:</u> **Dynamically splice specialized controllers (metaRL)**

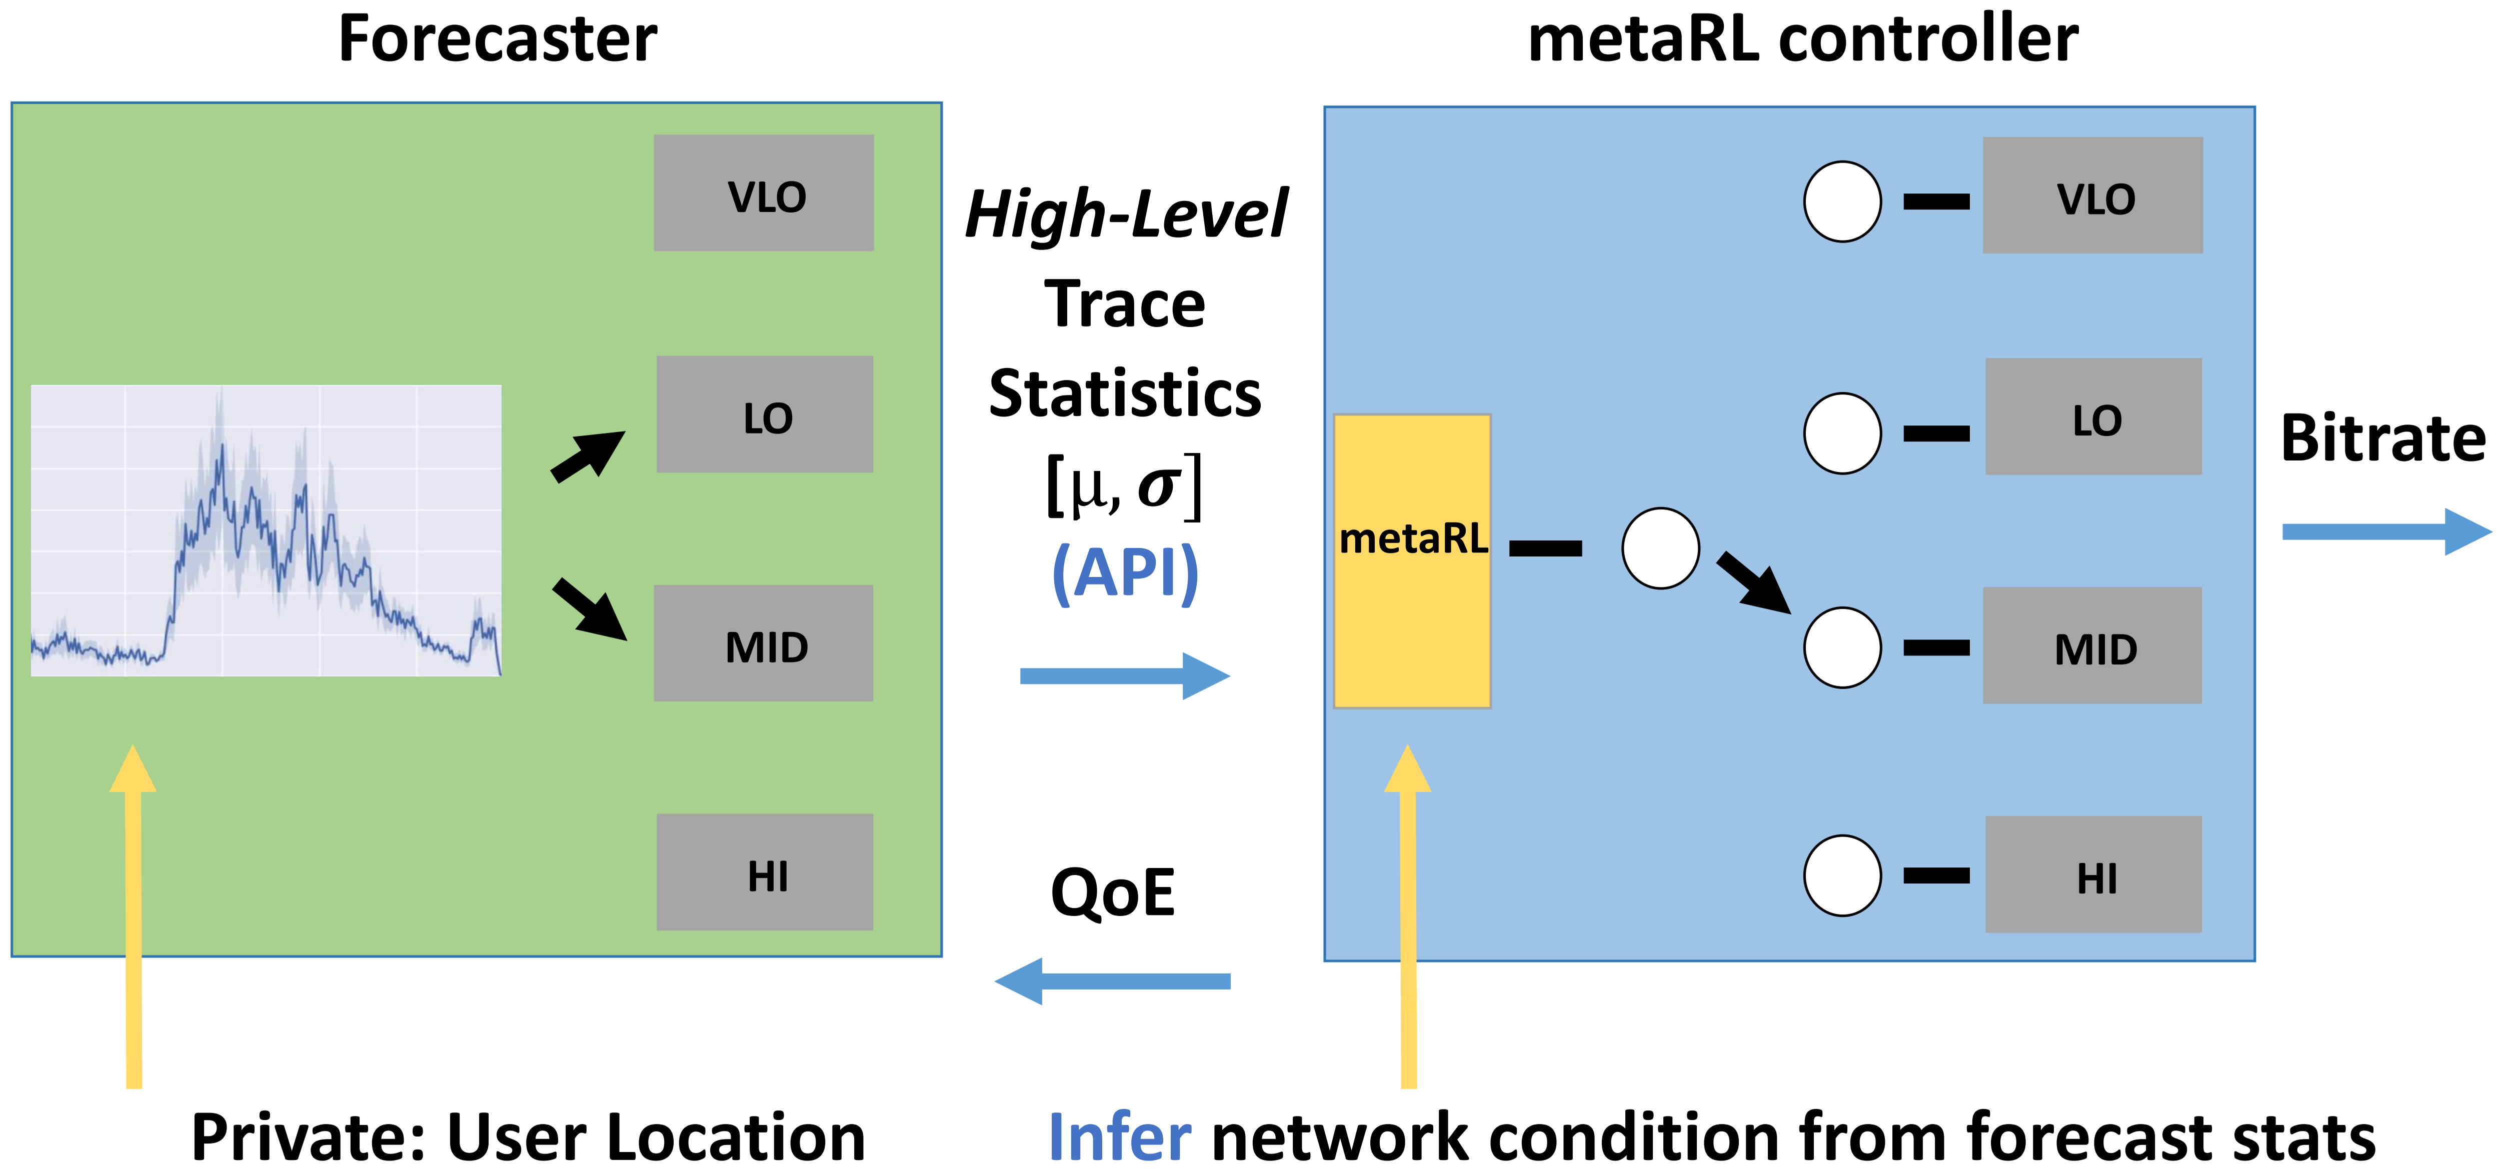
Forecaster
metaRL controller
VLO
LO
MID
HI
High-Level
Trace
Statistics
[μ, σ]
(API)
metaRL
VLO
LO
MID
HI
Bitrate
QoE
Private: User Location
Infer network condition from forecast stats

Palo Alto (Our data) + FCC/Norway (Pensieve)

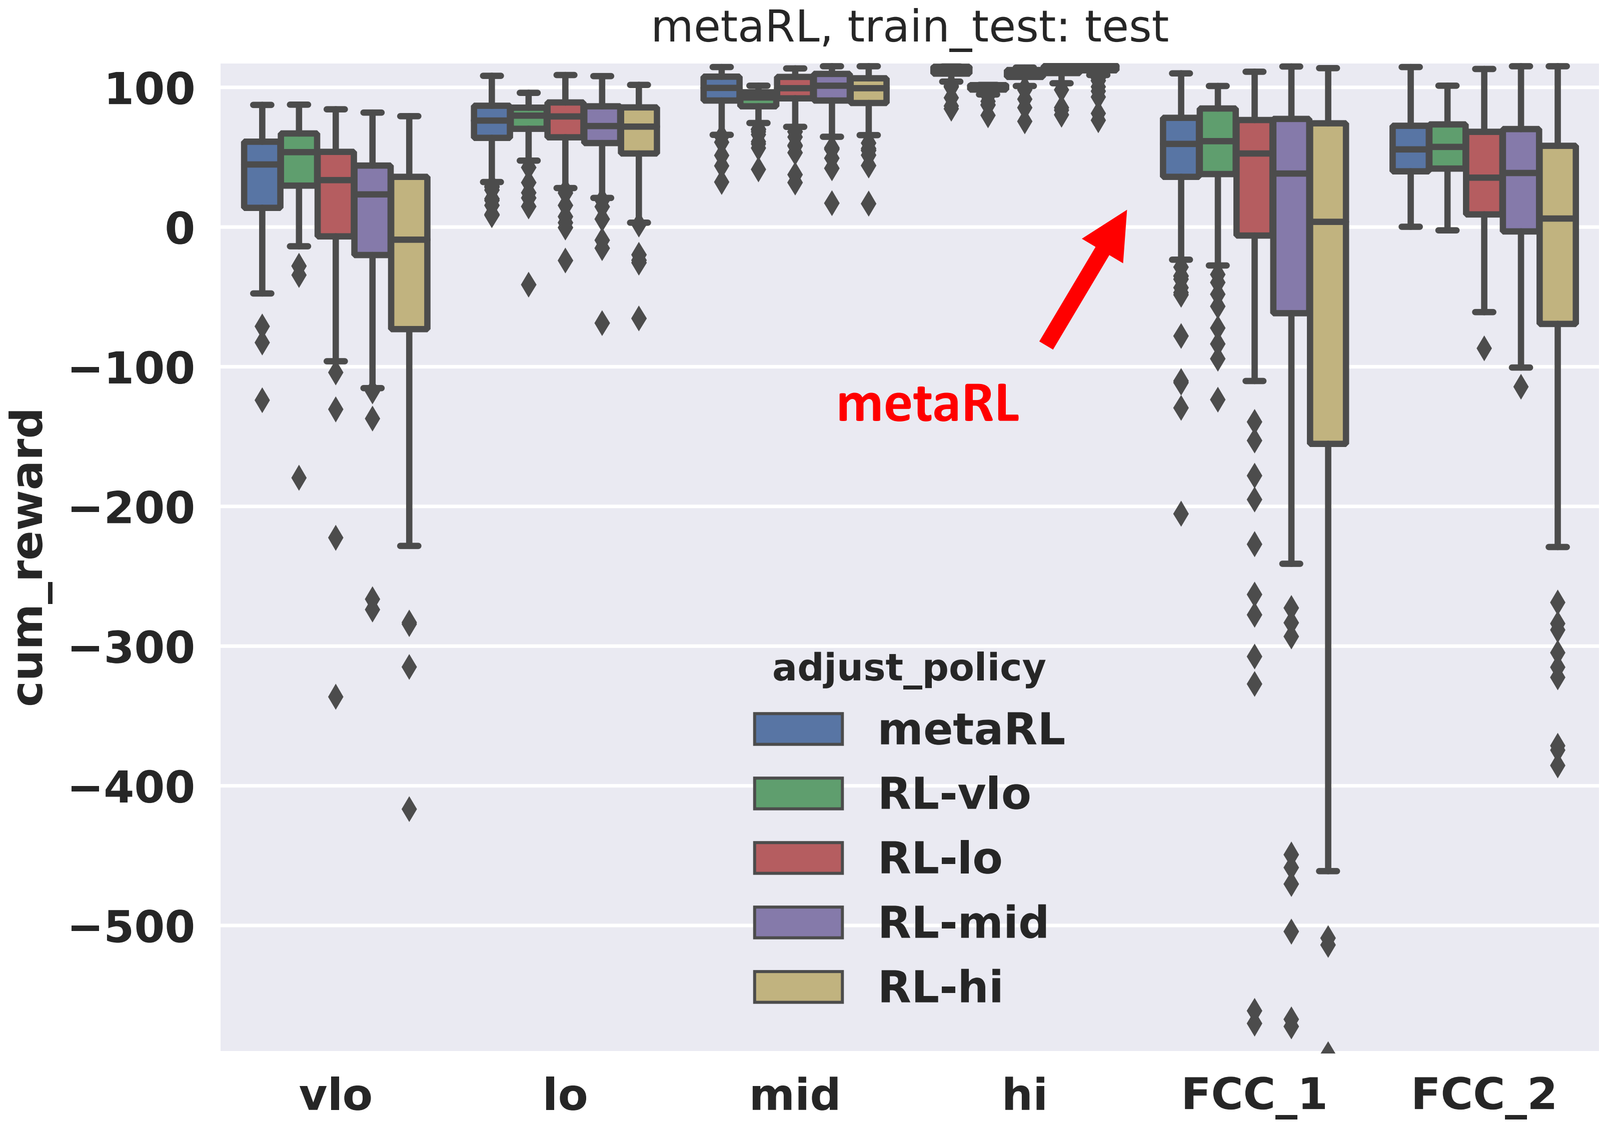

Generalize to FCC/Norway data from Pensieve

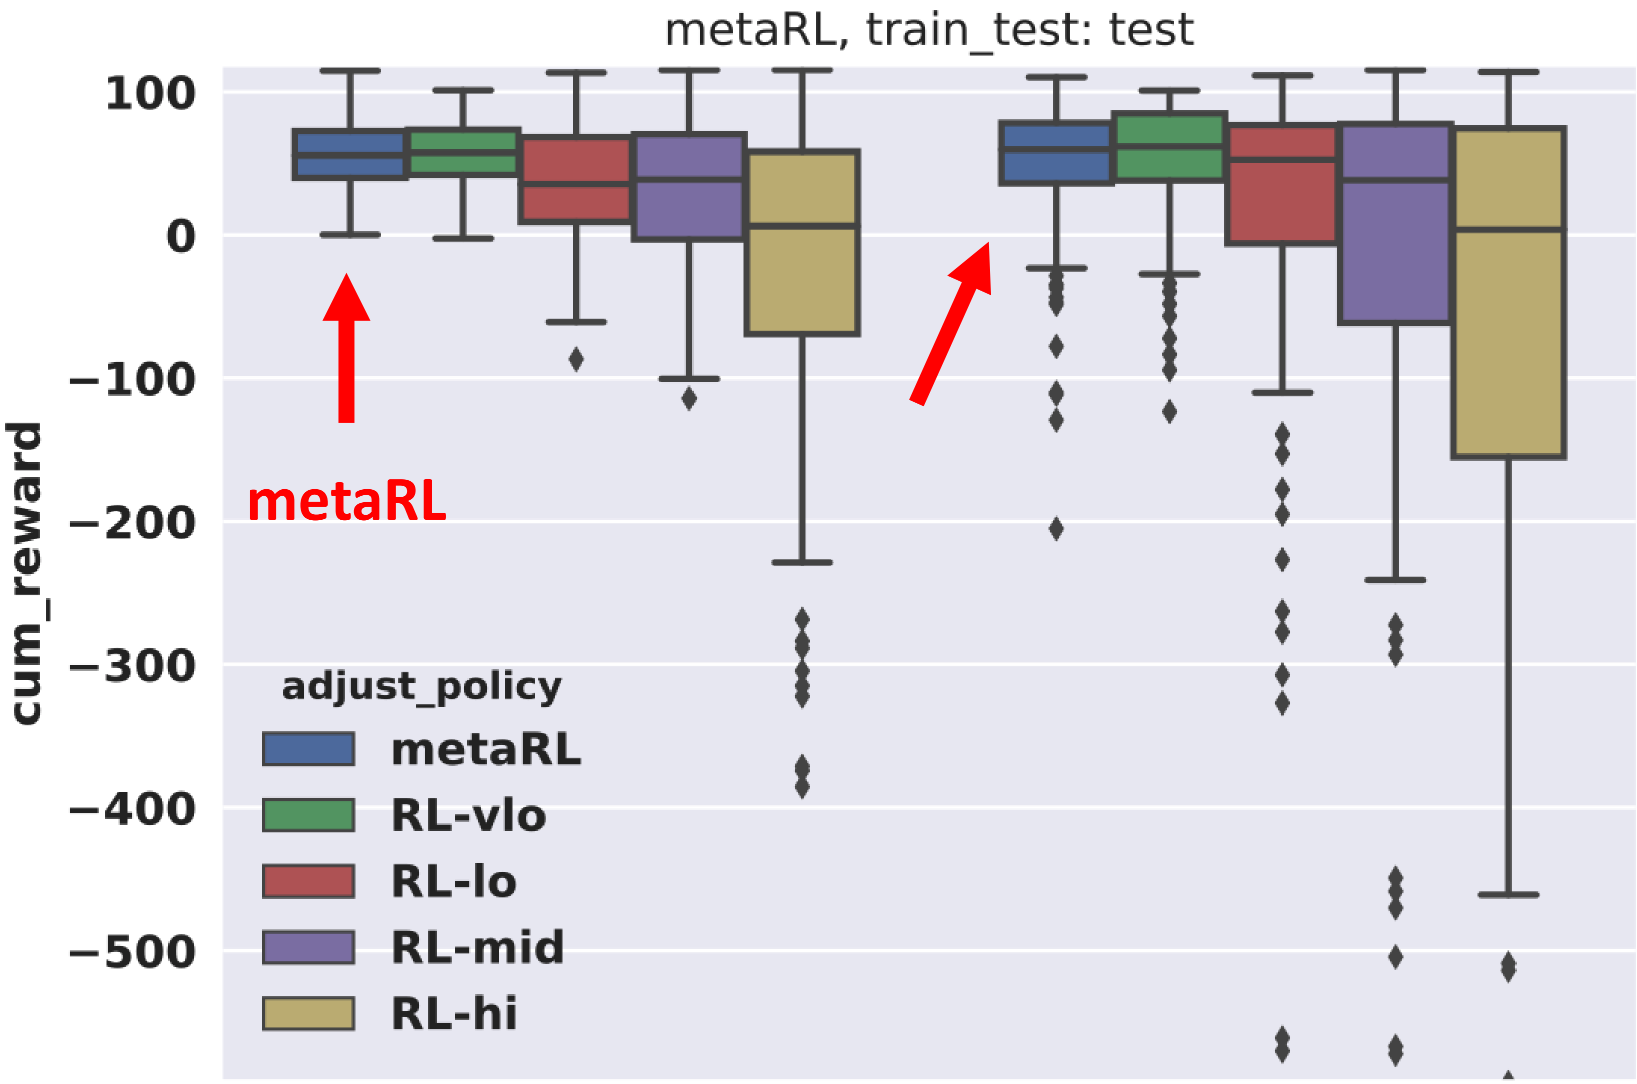

Re-analysis of Pensieve (Sigcomm 18, Mao et. al.)

$$\text{QoE} = \sum_{k=0}^{K} \text{Quality}(\text{Bitrate}) - \sum_{k=0}^{K} \text{Stalls} - \sum_{k=1}^{K-1} |\text{Quality}_{k+1} - \text{Quality}_k|$$

Linear QoE (hi-thpt)

HD QoE (vlo-thpt)

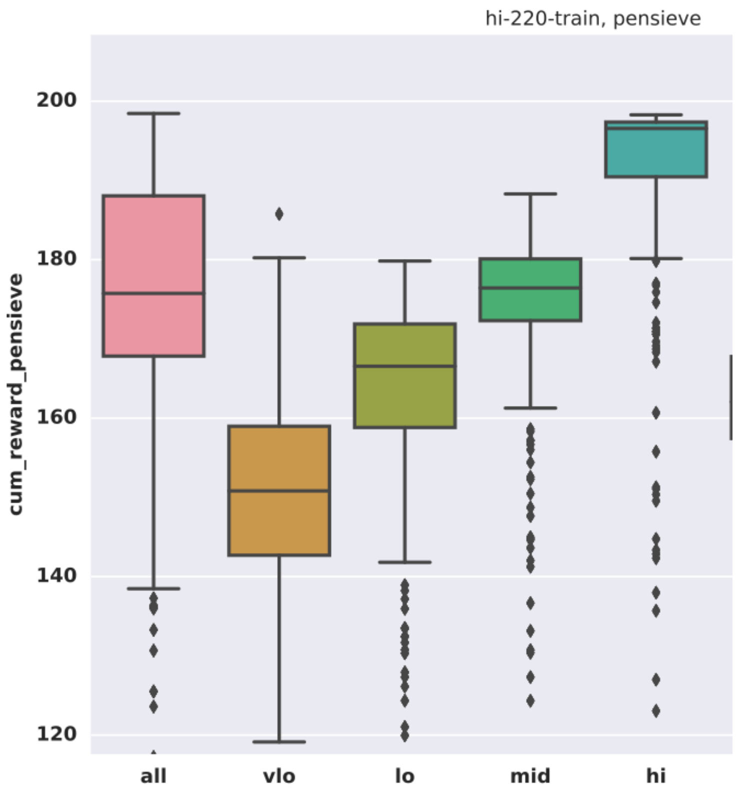

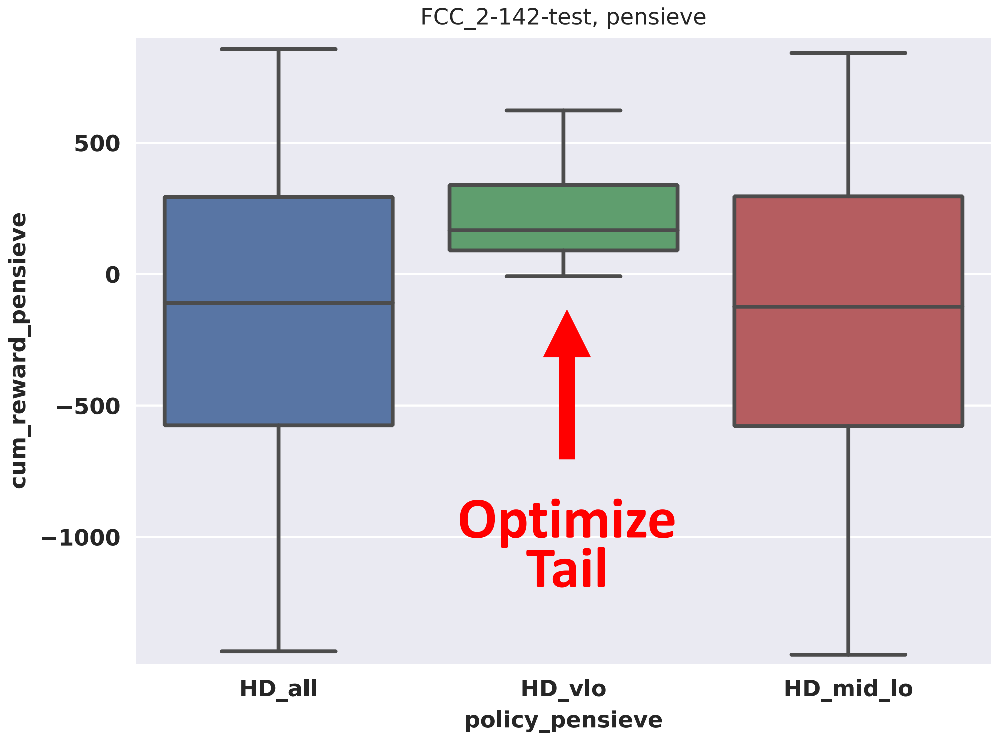

Future work

1. Broad-vision for Time-Series Control
   - Data-driven forecasts/ control strategies
   - ***Intrinsic data boundaries***

2. Value/Price of Information used for Long-Term Control?

3. Privacy/Information Leakage

Questions: csandeep@stanford.edu

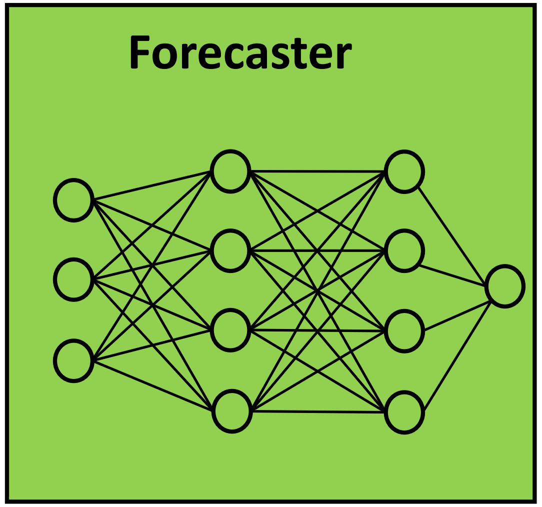

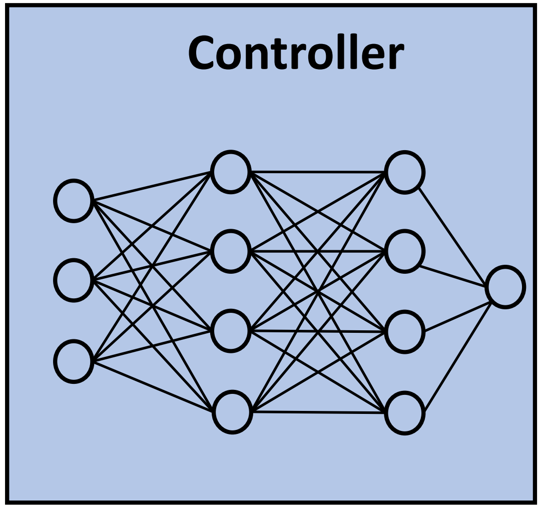

Extra slides

Claim: Decouple but *co-design* predictor and controller

Why not ***end-to-end*** learning?

Why **Decouple**?

1. Natural Data Boundaries
2. Modularity (Re-use forecaster)

Why **Co-design**?

1. Tune forecasts to control risk
2. Robust Adversarial Training

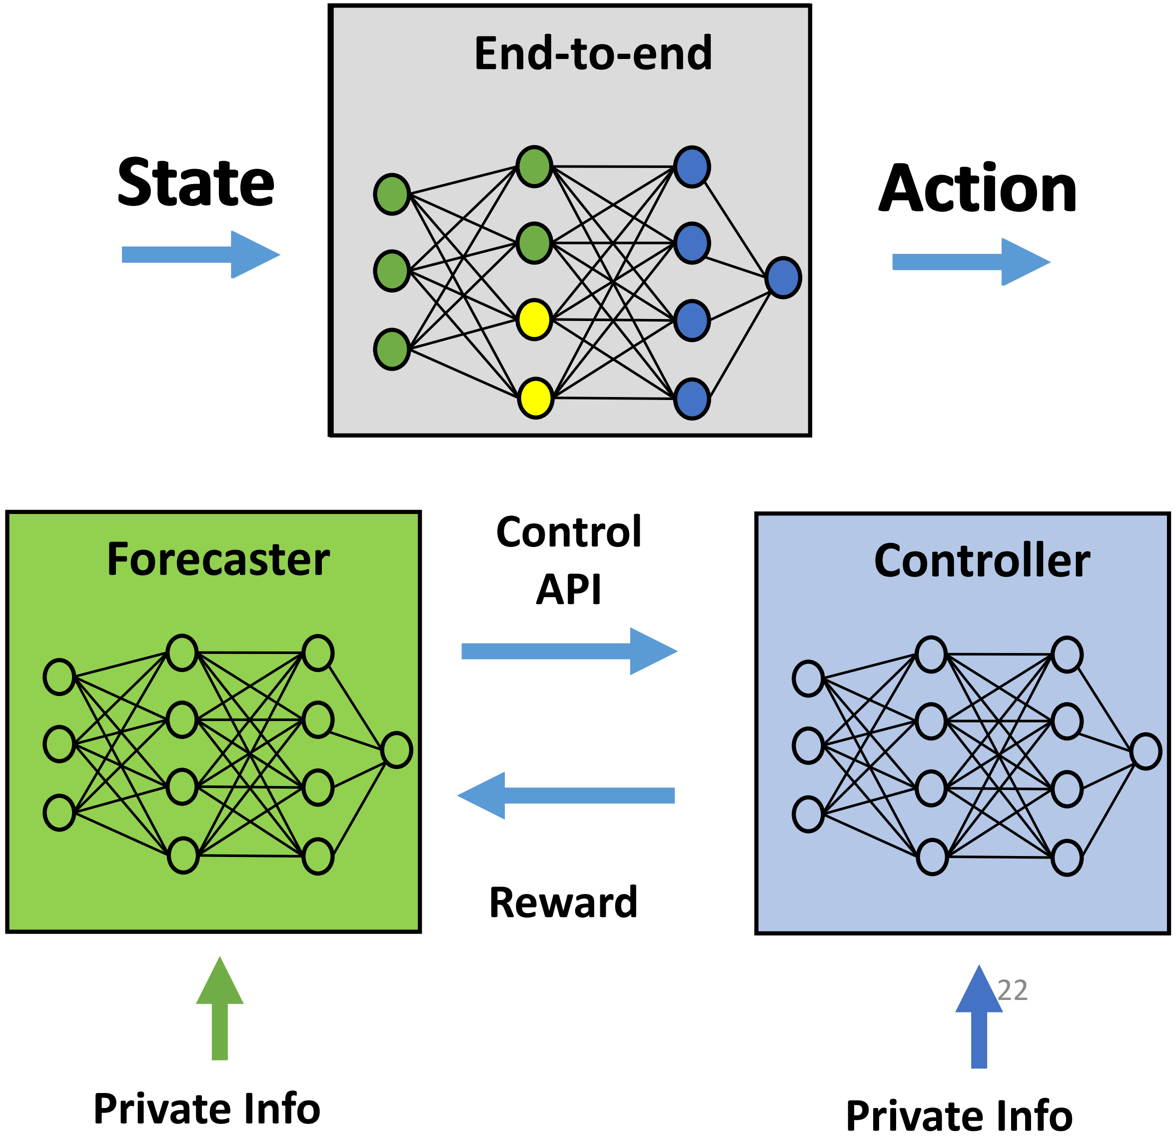

RL Formulation

$$\mathcal{M}^F = (S^F, A^F, \mathcal{T}^F, R^F)$$

$$\mathcal{M}^C = (S^C, A^C, \mathcal{T}^C, R^C)$$

Forecaster

$$a_t^F = \phi(s_t^F)$$

Controller

$$a_t^C$$

$$r_t^F = -r_t^C$$

Adversarial

$$s_t^F = \begin{bmatrix} x_t^{F,p} \\ x_t^J \end{bmatrix}$$

$$s_t^C = \begin{bmatrix} x_t^{C,p} \\ x_t^J \\ \phi(s_t^F) \end{bmatrix}$$

Quantifying Sub-optimality Gap

With oracle knowledge of network condition

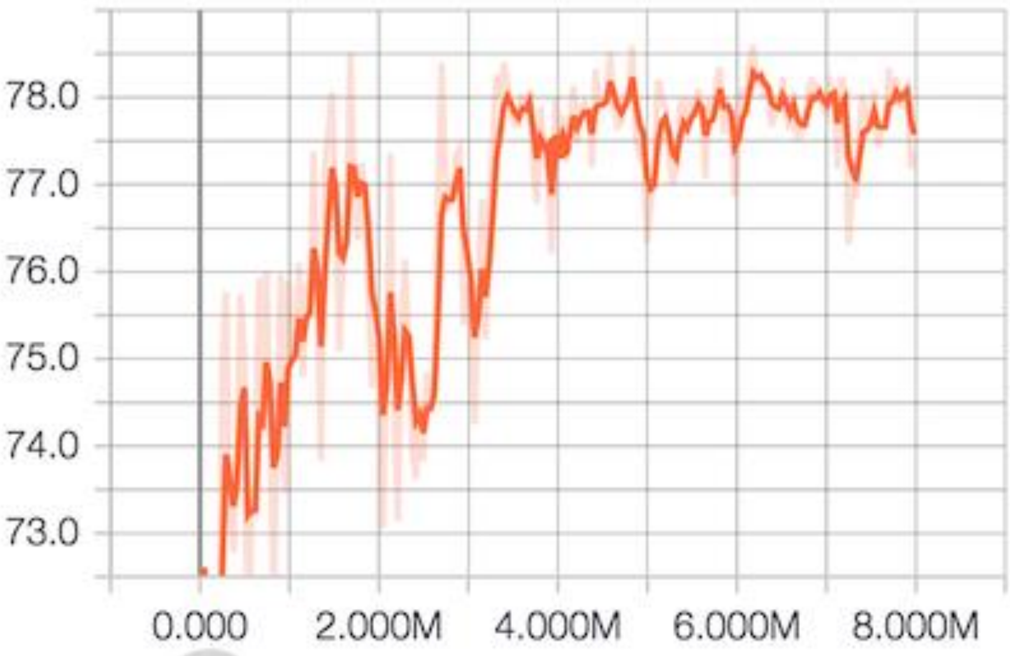

Gap

Have to learn network condition

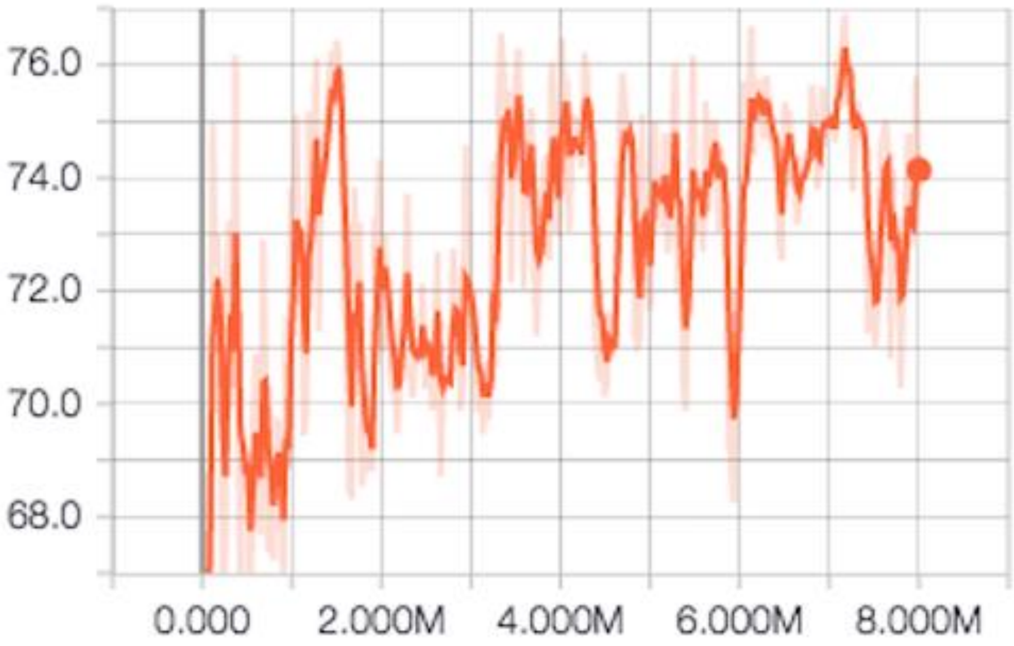

Value/Price of Timeseries Variables?

IoT Traffic Scheduling (AAAI 2018)

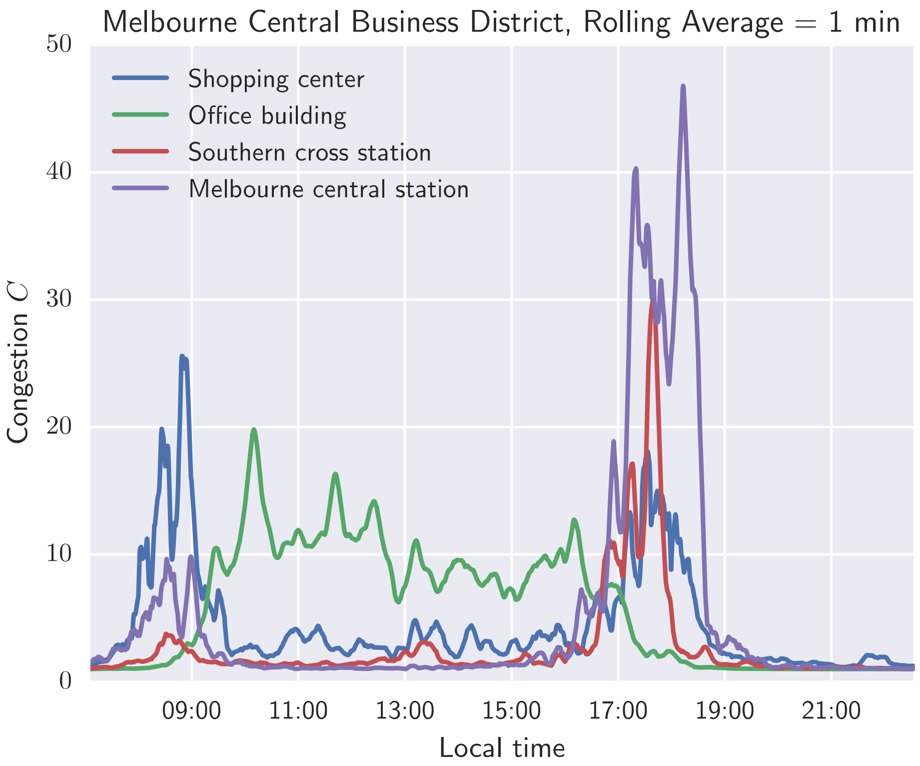

Diverse city-wide cell patterns

Network State

IoT Scheduler

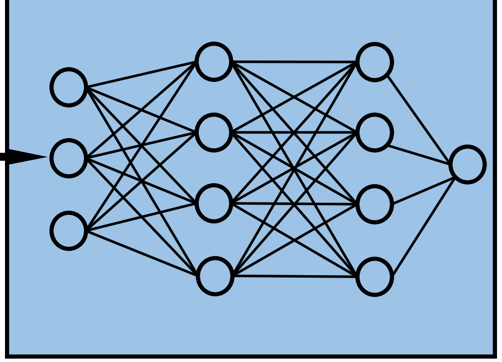

IoT rate